



ProCurve Series 6120 Switches

Access Security Guide

August 2009

© Copyright 2009 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. All Rights Reserved.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard.

Publication Number

5992-5525
August 2009

Applicable Products

HP ProCurve Switch 6120G/XG	(498358-B21)
HP ProCurve Switch 6120XG	(516733-B21)

Trademark Credits

Microsoft, Windows, and Microsoft Windows NT are U.S. registered trademarks of Microsoft Corporation.

Software Credits and Notices

SSH on ProCurve switches is based on the OpenSSH software toolkit. This product includes software developed by the OpenSSH Project for use in the OpenSSH Toolkit. For more information on OpenSSH, visit www.openssh.com.

SSL on ProCurve switches is based on the OpenSSL software toolkit. This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. For more information on OpenSSL, visit www.openssl.org.

This product includes cryptographic software written by Eric Young (ey@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Portions of the software on ProCurve switches are based on the lightweight TCP/IP (lwIP) software toolkit by Adam Dunkels, and are covered by the following notices.

Copyright © 2001-2003 Swedish Institute of Computer Science. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY,

OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes software written by Adam Dunkels (adam@sics.se).

Disclaimer

The information contained in this document is subject to change without notice.

HEWLETT-PACKARD COMPANY MAKES NO WARRANTY OF ANY KIND WITH REGARD TO THIS MATERIAL, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Hewlett-Packard assumes no responsibility for the use or reliability of its software on equipment that is not furnished by Hewlett-Packard.

Warranty

See the Customer Support/Warranty information at <http://www.hp.com/#Support>.

A copy of the specific warranty terms applicable to your Hewlett-Packard products and replacement parts can be obtained from your HP Sales and Service Office or authorized dealer.

HP Contact Information

For the name of the nearest HP authorized reseller:

- See the Contact HP worldwide (in English) webpage (<http://welcome.hp.com/country/us/en/wwwcontact.html>).

For HP technical support:

- In the United States, for contact options see the Contact HP United States webpage (http://welcome.hp.com/country/us/en/contact_us.html). To contact HP by phone:
 - Call 1-800-HP-INVENT (1-800-474-6836). This service is available 24 hours a day, 7 days a week. For continuous quality improvement, calls may be recorded or monitored.
 - If you have purchased a Care Pack (service upgrade), call 1-800-633-3600. For more information about Care Packs, refer to the HP website (<http://www.hp.com/hps>).
- In other locations, see the Contact HP worldwide (in English) webpage (<http://welcome.hp.com/country/us/en/wwwcontact.html>).

Contents

Product Documentation

About Your Switch Manual Set	xvii
Printed Publications.	xvii
Electronic Publications.	xvii
Software Feature Index	xviii

1 Security Overview

Contents	1-1
Introduction	1-2
About This Guide	1-2
For More Information	1-2
Access Security Features	1-3
Network Security Features	1-7
Getting Started with Access Security	1-9
Physical Security	1-9
Quick Start: Using the Management Interface Wizard	1-10
CLI: Management Interface Wizard	1-10
Web: Management Interface Wizard	1-12
SNMP Security Guidelines	1-15
Precedence of Security Options	1-17
Precedence of Port-Based Security Options	1-17
Precedence of Client-Based Authentication:	
Dynamic Configuration Arbiter	1-17
Network Immunity Manager	1-18
Arbitrating Client-Specific Attributes	1-19
ProCurve Identity-Driven Manager (IDM)	1-21

2 Configuring Username and Password Security

Contents	2-1
Overview	2-3
Configuring Local Password Security	2-6
Menu: Setting Passwords	2-6
CLI: Setting Passwords and Usernames	2-8
Web: Setting Passwords and Usernames	2-9
SNMP: Setting Passwords and Usernames	2-9
Saving Security Credentials in a Config File	2-10
Benefits of Saving Security Credentials	2-10
Enabling the Storage and Display of Security Credentials	2-11
Security Settings that Can Be Saved	2-11
Local Manager and Operator Passwords	2-12
Password Command Options	2-12
SNMP Security Credentials	2-13
802.1X Port-Access Credentials	2-14
TACACS+ Encryption Key Authentication	2-15
RADIUS Shared-Secret Key Authentication	2-15
SSH Client Public-Key Authentication	2-16
Operating Notes	2-19
Restrictions	2-21
Front-Panel Security	2-23
When Security Is Important	2-23
Front-Panel Button Functions	2-24
Clear Button	2-25
Reset Button	2-25
Restoring the Factory Default Configuration	2-25
Configuring Front-Panel Security	2-27
Disabling the Clear Password Function of the Clear Button ...	2-29
Re-Enabling the Clear Button and Setting or	
Changing the “Reset-On-Clear” Operation	2-30
Changing the Operation of the Reset+Clear Combination	2-31
Password Recovery	2-32
Disabling or Re-Enabling the Password Recovery Process	2-32

Password Recovery Process	2-34
---------------------------------	------

3 Web and MAC Authentication

Contents	3-1
Overview	3-3
Web Authentication	3-3
MAC Authentication	3-4
Concurrent Web and MAC Authentication	3-4
Authorized and Unauthorized Client VLANs	3-5
RADIUS-Based Authentication	3-6
Wireless Clients	3-6
How Web and MAC Authentication Operate	3-6
Web-based Authentication	3-7
MAC-based Authentication	3-9
Terminology	3-11
Operating Rules and Notes	3-12
Setup Procedure for Web/MAC Authentication	3-14
Before You Configure Web/MAC Authentication	3-14
Configuring the RADIUS Server To Support MAC Authentication ..	3-16
Configuring the Switch To Access a RADIUS Server	3-17
Configuring Web Authentication	3-20
Overview	3-20
Configuration Commands for Web Authentication	3-21
Show Commands for Web Authentication	3-28
Customizing Web Authentication HTML Files (Optional)	3-34
Implementing Customized Web-Auth Pages	3-34
Operating Notes and Guidelines	3-34
Customizing HTML Templates	3-35
Customizable HTML Templates	3-36
Configuring MAC Authentication on the Switch	3-50
Overview	3-50
Configuration Commands for MAC Authentication	3-51
Show Commands for MAC-Based Authentication	3-54

Client Status	3-60
----------------------------	------

4 TACACS+ Authentication

Contents	4-1
Overview	4-2
Terminology Used in TACACS Applications:	4-3
General System Requirements	4-5
General Authentication Setup Procedure	4-5
Configuring TACACS+ on the Switch	4-8
Before You Begin	4-8
CLI Commands Described in this Section	4-9
Viewing the Switch's Current Authentication Configuration	4-9
Viewing the Switch's Current TACACS+ Server Contact Configuration	4-10
Configuring the Switch's Authentication Methods	4-11
Using the Privilege-Mode Option for Login	4-11
Authentication Parameters	4-12
Configuring the TACACS+ Server for Single Login	4-13
Configuring the Switch's TACACS+ Server Access	4-18
How Authentication Operates	4-24
General Authentication Process Using a TACACS+ Server	4-24
Local Authentication Process	4-26
Using the Encryption Key	4-27
General Operation	4-27
Encryption Options in the Switch	4-27
Controlling Web Browser Interface	
Access When Using TACACS+ Authentication	4-28
Messages Related to TACACS+ Operation	4-29
Operating Notes	4-29

5 RADIUS Authentication, Authorization, and Accounting

Contents	5-1
Overview	5-3

Authentication Services	5-3
Accounting Services	5-4
RADIUS-Administered CoS and Rate-Limiting	5-4
RADIUS-Administered Commands Authorization	5-4
SNMP Access to the Switch's Authentication Configuration MIB ...	5-4
Terminology	5-5
Switch Operating Rules for RADIUS	5-6
General RADIUS Setup Procedure	5-7
Configuring the Switch for RADIUS Authentication	5-8
Outline of the Steps for Configuring RADIUS Authentication	5-9
1. Configure Authentication for the Access Methods You Want RADIUS To Protect	5-10
2. Enable the (Optional) Access Privilege Option	5-13
3. Configure the Switch To Access a RADIUS Server	5-14
4. Configure the Switch's Global RADIUS Parameters	5-17
Using SNMP To View and Configure	
Switch Authentication Features	5-21
Changing and Viewing the SNMP Access Configuration	5-22
Local Authentication Process	5-24
Controlling Web Browser Interface Access	5-25
Commands Authorization	5-26
Enabling Authorization	5-27
Displaying Authorization Information	5-28
Configuring Commands Authorization on a RADIUS Server	5-28
Using Vendor Specific Attributes (VSAs)	5-28
Example Configuration on Cisco Secure ACS for MS Windows	5-30
Example Configuration Using FreeRADIUS	5-32
VLAN Assignment in an Authentication Session	5-34
Tagged and Untagged VLAN Attributes	5-35
Additional RADIUS Attributes	5-36
Configuring RADIUS Accounting	5-37
Operating Rules for RADIUS Accounting	5-39
Steps for Configuring RADIUS Accounting	5-39
1. Configure the Switch To Access a RADIUS Server	5-40

2. Configure Accounting Types and the Controls for Sending Reports to the RADIUS Server	5-42
3. (Optional) Configure Session Blocking and Interim Updating Options	5-44
Viewing RADIUS Statistics	5-46
General RADIUS Statistics	5-46
RADIUS Authentication Statistics	5-48
RADIUS Accounting Statistics	5-49
Changing RADIUS-Server Access Order	5-50
Messages Related to RADIUS Operation	5-53

6 Configuring Secure Shell (SSH)

Contents	6-1
Overview	6-2
Terminology	6-3
Prerequisite for Using SSH	6-5
Public Key Formats	6-5
Steps for Configuring and Using SSH for Switch and Client Authentication	6-6
General Operating Rules and Notes	6-8
Configuring the Switch for SSH Operation	6-9
1. Assigning a Local Login (Operator) and Enable (Manager) Password	6-10
2. Generating the Switch's Public and Private Key Pair	6-10
Configuring Key Lengths	6-13
3. Providing the Switch's Public Key to Clients	6-13
4. Enabling SSH on the Switch and Anticipating SSH Client Contact Behavior	6-15
5. Configuring the Switch for SSH Authentication	6-20
6. Use an SSH Client To Access the Switch	6-24
Further Information on SSH Client Public-Key Authentication .	6-24
Messages Related to SSH Operation	6-30
Logging Messages	6-31
Debug Logging	6-32

7 Configuring Secure Socket Layer (SSL)

Contents	7-1
Overview	7-2
Terminology	7-3
Prerequisite for Using SSL	7-5
Steps for Configuring and Using SSL for Switch and Client Authentication	7-5
General Operating Rules and Notes	7-6
Configuring the Switch for SSL Operation	7-7
1. Assigning a Local Login (Operator) and Enabling (Manager) Password	7-7
2. Generating the Switch's Server Host Certificate	7-8
To Generate or Erase the Switch's Server Certificate with the CLI	7-9
Comments on Certificate Fields.	7-10
Generate a Self-Signed Host Certificate with the Web browser interface	7-12
Generate a CA-Signed server host certificate with the Web browser interface	7-15
3. Enabling SSL on the Switch and Anticipating SSL Browser Contact Behavior	7-17
Using the CLI Interface to Enable SSL	7-19
Using the Web Browser Interface to Enable SSL	7-19
Common Errors in SSL setup	7-21

8 Configuring Advanced Threat Protection

Contents	8-1
Introduction	8-3
DHCP Snooping	8-4
Overview	8-4
Enabling DHCP Snooping	8-5
Enabling DHCP Snooping on VLANs	8-7
Configuring DHCP Snooping Trusted Ports	8-8
Configuring Authorized Server Addresses	8-9

Using DHCP Snooping with Option 82	8-9
Changing the Remote-id from a MAC to an IP Address	8-11
Disabling the MAC Address Check	8-11
The DHCP Binding Database	8-12
Operational Notes	8-13
Log Messages	8-14
Dynamic ARP Protection	8-16
Introduction	8-16
Enabling Dynamic ARP Protection	8-18
Configuring Trusted Ports	8-18
Adding an IP-to-MAC Binding to the DHCP Database	8-20
Configuring Additional Validation Checks on ARP Packets	8-21
Verifying the Configuration of Dynamic ARP Protection	8-21
Displaying ARP Packet Statistics	8-22
Monitoring Dynamic ARP Protection	8-23
Dynamic IP Lockdown	8-23
Protection Against IP Source Address Spoofing	8-24
Prerequisite: DHCP Snooping	8-24
Filtering IP and MAC Addresses Per-Port and Per-VLAN	8-25
Enabling Dynamic IP Lockdown	8-26
Operating Notes	8-26
Adding an IP-to-MAC Binding to the DHCP Binding Database	8-28
Potential Issues with Bindings	8-28
Adding a Static Binding	8-29
Verifying the Dynamic IP Lockdown Configuration	8-29
Displaying the Static Configuration of IP-to-MAC Bindings	8-30
Debugging Dynamic IP Lockdown	8-31
Using the Instrumentation Monitor	8-33
Operating Notes	8-34
Configuring Instrumentation Monitor	8-35
Examples	8-36
Viewing the Current Instrumentation Monitor Configuration	8-37

9 Traffic/Security Filters and Monitors

Contents	9-1
Overview	9-2
Introduction	9-2
Filter Limits	9-2
Using Port Trunks with Filters	9-2
Filter Types and Operation	9-3
Source-Port Filters	9-4
Operating Rules for Source-Port Filters	9-4
Example	9-5
Named Source-Port Filters	9-6
Operating Rules for Named Source-Port Filters	9-6
Defining and Configuring Named Source-Port Filters	9-7
Viewing a Named Source-Port Filter	9-8
Using Named Source-Port Filters	9-9
Configuring Traffic/Security Filters	9-15
Configuring a Source-Port Traffic Filter	9-16
Example of Creating a Source-Port Filter	9-17
Configuring a Filter on a Port Trunk	9-17
Editing a Source-Port Filter	9-18
Filter Indexing	9-19
Displaying Traffic/Security Filters	9-20

10 Configuring Port-Based and User-Based Access Control (802.1X)

Contents	10-1
Overview	10-3
Why Use Port-Based or User-Based Access Control?	10-3
General Features	10-3
User Authentication Methods	10-4
802.1X User-Based Access Control	10-4
802.1X Port-Based Access Control	10-5
Alternative To Using a RADIUS Server	10-6
Accounting	10-6

Terminology	10-6
General 802.1X Authenticator Operation	10-9
Example of the Authentication Process	10-9
VLAN Membership Priority	10-10
General Operating Rules and Notes	10-12
General Setup Procedure for 802.1X Access Control	10-14
Do These Steps Before You Configure 802.1X Operation	10-14
Overview: Configuring 802.1X Authentication on the Switch	10-17
Configuring Switch Ports as 802.1X Authenticators	10-18
1. Enable 802.1X Authentication on Selected Ports	10-19
A. Enable the Selected Ports as Authenticators and Enable the (Default) Port-Based Authentication	10-19
B. Specify User-Based Authentication or Return to Port-Based Authentication	10-20
Example: Configuring User-Based 802.1X Authentication	10-21
Example: Configuring Port-Based 802.1X Authentication	10-21
2. Reconfigure Settings for Port-Access	10-21
3. Configure the 802.1X Authentication Method	10-24
4. Enter the RADIUS Host IP Address(es)	10-25
5. Enable 802.1X Authentication on the Switch	10-25
6. Optional: Reset Authenticator Operation	10-26
7. Optional: Configure 802.1X Controlled Directions	10-26
Wake-on-LAN Traffic	10-27
Operating Notes	10-27
Example: Configuring 802.1X Controlled Directions	10-28
802.1X Open VLAN Mode	10-29
Introduction	10-29
VLAN Membership Priorities	10-30
Use Models for 802.1X Open VLAN Modes	10-31
Operating Rules for Authorized-Client and Unauthorized-Client VLANs	10-36
Setting Up and Configuring 802.1X Open VLAN Mode	10-40
802.1X Open VLAN Operating Notes	10-44
Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1X-Authenticated Devices	10-45

Port-Security	10-46
Configuring Switch Ports To Operate As Supplicants for 802.1X	
Connections to Other Switches	10-47
Example	10-47
Supplicant Port Configuration	10-49
Displaying 802.1X Configuration, Statistics, and Counters	10-51
Show Commands for Port-Access Authenticator	10-51
Viewing 802.1X Open VLAN Mode Status	10-60
Show Commands for Port-Access Supplicant	10-64
How RADIUS/802.1X Authentication Affects VLAN Operation .	10-65
VLAN Assignment on a Port	10-66
Operating Notes	10-66
Example of Untagged VLAN Assignment in a RADIUS-Based	
Authentication Session	10-68
Enabling the Use of GVRP-Learned Dynamic VLANs	
in Authentication Sessions	10-71
Messages Related to 802.1X Operation	10-73

11 Configuring and Monitoring Port Security

Contents	11-1
Overview	11-3
Port Security	11-4
Basic Operation	11-4
Eavesdrop Protection	11-5
Blocking Unauthorized Traffic	11-5
Trunk Group Exclusion	11-6
Planning Port Security	11-7
Port Security Command Options and Operation	11-8
Port Security Display Options	11-8
Configuring Port Security	11-12
Retention of Static Addresses	11-17
MAC Lockdown	11-22
Differences Between MAC Lockdown and Port Security	11-24
MAC Lockdown Operating Notes	11-25

Deploying MAC Lockdown	11-26
MAC Lockout	11-26
Port Security and MAC Lockout	11-29
Web: Displaying and Configuring Port Security Features	11-30
Reading Intrusion Alerts and Resetting Alert Flags	11-30
Notice of Security Violations	11-30
How the Intrusion Log Operates	11-31
Keeping the Intrusion Log Current by Resetting Alert Flags	11-32
Menu: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags	11-33
CLI: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags	11-34
Using the Event Log To Find Intrusion Alerts	11-36
Web: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags	11-37
Operating Notes for Port Security	11-38

12 Using Authorized IP Managers

Contents	12-1
Overview	12-2
Options	12-3
Access Levels	12-3
Defining Authorized Management Stations	12-4
Overview of IP Mask Operation	12-4
Menu: Viewing and Configuring IP Authorized Managers	12-5
CLI: Viewing and Configuring Authorized IP Managers	12-6
Listing the Switch's Current Authorized IP Manager(s)	12-6
Configuring IP Authorized Managers for the Switch	12-7
Web: Configuring IP Authorized Managers	12-9
Web Proxy Servers	12-9
How to Eliminate the Web Proxy Server	12-9
Using a Web Proxy Server to Access the Web Browser Interface ... 12-10	
Web-Based Help	12-10

Building IP Masks	12-10
Configuring One Station Per Authorized Manager IP Entry	12-10
Configuring Multiple Stations Per Authorized Manager IP Entry ..	12-11
Additional Examples for Authorizing Multiple Stations	12-13
Operating Notes	12-13

Product Documentation

About Your Switch Manual Set

Note

For the latest version of switch documentation, please visit any of the following websites:

www.procurve.com/manuals

www.hp.com/go/bladesystem/documentation

h18004.www1.hp.com/products/blades/components/c-class-tech-installing.html

Printed Publications

The publication listed below is printed and shipped with your switch. The latest version is also available in PDF format, as described in the Note at the top of this page.

- *Quick Installation Card*—Explains how to prepare for and perform the physical installation and connect the switch to your network.

Electronic Publications

The latest version of each of the publications listed below is available in PDF format on the ProCurve Web site, as described in the Note at the top of this page.

- *Installation and Getting Started Guide*—Describes how to install hardware, and get started using the switch.
- *Management and Configuration Guide*—Describes how to configure, manage, and monitor basic switch operation.
- *Advanced Traffic Management Guide*—Explains how to configure traffic management features such as VLANs, MSTP, and QoS.
- *Multicast and Routing Guide*—Explains how to configure IGMP features.
- *Access Security Guide*—Explains how to configure access security features and user authentication on the switch.
- *IPv6 Configuration Guide*—Describes the IPv6 protocol operations that are supported on the switch.
- *Release Notes*—Describe new features, fixes, and enhancements that become available between revisions of the main product guide.

Software Feature Index

This feature index indicates which manual to consult for information on a given software feature.

Note

This Index does not cover IPv6 capable software features. For information on IPv6 protocol operations and features (such as DHCPv6, DNS for IPv6, and Ping6), refer to the *IPv6 Configuration Guide*.

Intelligent Edge Software Features. These features are automatically included on all switches.

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
802.1Q VLAN Tagging		X		
802.1X Multiple Authenticated Clients Per Port				X
AAA Authentication				X
Authorized IP Managers				X
Authorized Manager List (Web, Telnet, TFTP)				X
Auto MDIX Configuration	X			
BOOTP	X			
Config File	X			
Console Access	X			
Copy Command	X			
CoS (Class of Service)		X		
Debug	X			
DHCP Configuration	X			
DHCP/Bootp Operation	X			
Diagnostic Tools	X			

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
Downloading Software	X			
Event Log	X			
Factory Default Settings	X			
Flow Control (802.3x)	X			
File Transfers	X			
Friendly Port Names	X			
GVRP		X		
Identity-Driven Management (IDM)		X		
IGMP			X	
Interface Access (Telnet, Console/Serial, Web)	X			
IP Addressing	X			
Jumbo Packets	X			
LACP	X			
LLDP	X			
LLDP-MED	X			
Loop Protection		X		
MAC Address Management	X			
MAC Lockdown				X
MAC Lockout				X
MAC-based Authentication				X
Monitoring and Analysis	X			
Network Management Applications (SNMP)	X			
Passwords and Password Clear Protection				X
ProCurve Manager (PCM)	X			
Ping	X			
Port Configuration	X			

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
Port Monitoring		X		
Port Security				X
Port Status	X			
Port Trunking (LACP)	X			
Port-Based Access Control (802.1X)				X
Protocol VLANs		X		
Quality of Service (QoS)		X		
RADIUS Authentication and Accounting				X
RADIUS-Based Configuration				X
RMON 1,2,3,9	X			
Secure Copy	X			
SFTP	X			
SNMPv3	X			
Software Downloads (SCP/SFTP, TFTP, Xmodem)	X			
Source-Port Filters				X
Spanning Tree (STP, RSTP, MSTP)		X		
SSHv2 (Secure Shell) Encryption				X
SSL (Secure Socket Layer)				X
Syslog	X			
System Information	X			
TACACS+ Authentication				X
Telnet Access	X			
TFTP	X			
Time Protocols (TimeP, SNTP)	X			
Troubleshooting	X			
Uni-Directional Link Detection (UDLD)	X			

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
VLANs		X		
Web Authentication RADIUS Support				X
Web-based Authentication				X
Web UI	X			

Security Overview

Contents

- Introduction** 1-2
 - About This Guide 1-2
 - For More Information 1-2
- Access Security Features** 1-3
- Network Security Features** 1-7
- Getting Started with Access Security** 1-9
 - Physical Security 1-9
 - Quick Start: Using the Management Interface Wizard 1-10
 - CLI: Management Interface Wizard 1-10
 - Web: Management Interface Wizard 1-12
 - SNMP Security Guidelines 1-15
- Precedence of Security Options** 1-17
 - Precedence of Port-Based Security Options 1-17
 - Precedence of Client-Based Authentication:
 - Dynamic Configuration Arbiter 1-17
 - Network Immunity Manager 1-18
 - Arbitrating Client-Specific Attributes 1-19
- ProCurve Identity-Driven Manager (IDM)** 1-21

Introduction

This chapter provides an overview of the security features included on your switch. Table 1-1 on page 1-3 outlines the access security and authentication features, while Table 1-2 on page 1-7 highlights the additional features designed to help secure and protect your network. For detailed information on individual features, see the references provided.

Before you connect your switch to a network, ProCurve strongly recommends that you review the section titled “Getting Started with Access Security” on page 1-9. It outlines potential threats for unauthorized switch and network access, and provides guidelines on how to prepare the switch for secure network operation.

About This Guide

This *Access Security Guide* describes how to configure security features on your switch.

Note

For an introduction to the standard conventions used in this guide, refer to the *Getting Started* chapter in the *Management and Configuration Guide* for your switch.

For More Information

For IPv6-specific security settings and features, refer to the *IPv6 Configuration Guide* for your switch.

For information on which product manual to consult for a specific software feature, refer to the “Software Feature Index” on page xviii of this guide.

For the latest version of all HP ProCurve switch documentation, including Release Notes covering recently added features and other software topics, visit the following web sites:

www.procurve.com/manuals

www.hp.com/go/ bladesystem/documentation

h18004.www1.hp.com/products/blades/components/c-class-tech-installing.html

Access Security Features

This section provides an overview of the switch’s access security features, authentication protocols, and methods. Table 1-1 lists these features and provides summary configuration guidelines. For more in-depth information, see the references provided (all chapter and page references are to this *Access Security Guide* unless a different manual name is indicated).

Note

The Management Interface wizard provides a convenient step-by-step method to prepare the switch for secure network operation. See “Quick Start: Using the Management Interface Wizard” on page 1-10 for details.

Table 1-1. Access Security and Switch Authentication Features

Feature	Default Setting	Security Guidelines	More Information and Configuration Details
Manager password	no password	<p>Configuring a local Manager password is a fundamental step in reducing the possibility of unauthorized access through the switch’s Web browser and console (CLI and Menu) interfaces. The Manager password can easily be set by any one of the following methods:</p> <ul style="list-style-type: none">• CLI: password manager command, or Management interface wizard• Web browser interface: the password options under the Security tab, or Management interface wizard• Menu interface: Console Passwords option• SNMP	<p><i>“Configuring Local Password Security” on page 2-6</i></p> <p><i>“Quick Start: Using the Management Interface Wizard” on page 1-10</i></p> <p><i>“Using SNMP To View and Configure Switch Authentication Features” on page 6-21</i></p>

Feature	Default Setting	Security Guidelines	More Information and Configuration Details
Telnet and Web-browser access	enabled	<p>The default remote management protocols enabled on the switch are plain text protocols, which transfer passwords in open or plain text that is easily captured. To reduce the chances of unauthorized users capturing your passwords, secure and encrypted protocols such as SSH and SSL (see below for details) should be used for remote access. This enables you to employ increased access security while still retaining remote client access.</p> <p>Also, access security on the switch is incomplete without disabling Telnet and the standard Web browser access. Among the methods for blocking unauthorized access attempts using Telnet or the Web browser are the following two CLI commands:</p> <ul style="list-style-type: none"> • no telnet-server: This command blocks inbound Telnet access. • no web-management: This command prevents use of the Web browser interface through http (port 80) server access. <p>If you choose not to disable Telnet and Web browser access, you may want to consider using RADIUS accounting to maintain a record of password-protected access to the switch.</p>	<p><i>“Quick Start: Using the Management Interface Wizard” on page 1-10</i></p> <p>For more on Telnet and web browser access, refer to the chapter on <i>“Interface Access and System Information”</i> in the <i>Management and Configuration Guide</i>.</p> <p>For RADIUS accounting, refer to <i>Chapter 6, “RADIUS Authentication and Accounting”</i></p>
SSH	disabled	<p>SSH provides Telnet-like functions through encrypted, authenticated transactions of the following types:</p> <ul style="list-style-type: none"> • client public-key authentication: uses one or more public keys (from clients) that must be stored on the switch. Only a client with a private key that matches a stored public key can gain access to the switch. • switch SSH and user password authentication: this option is a subset of the client public-key authentication, and is used if the switch has SSH enabled without a login access configured to authenticate the client’s key. In this case, the switch authenticates itself to clients, and users on SSH clients then authenticate themselves to the switch by providing passwords stored on a RADIUS or TACACS+ server, or locally on the switch. • secure copy (SC) and secure FTP (SFTP): By opening a secure, encrypted SSH session, you can take advantage of SC and SFTP to provide a secure alternative to TFTP for transferring sensitive switch information. For more on SC and SFTP, refer to the section titled <i>“Using Secure Copy and SFTP”</i> in the <i>“File Transfers”</i> appendix of the <i>Management and Configuration Guide</i> for your switch. 	<p><i>“Quick Start: Using the Management Interface Wizard” on page 1-10</i></p> <p><i>Chapter 8 “Configuring Secure Shell (SSH)”</i></p>

Feature	Default Setting	Security Guidelines	More Information and Configuration Details
SSL	disabled	Secure Socket Layer (SSL) and Transport Layer Security (TLS) provide remote Web browser access to the switch via authenticated transactions and encrypted paths between the switch and management station clients capable of SSL/TLS operation. The authenticated type includes server certificate authentication with user password authentication.	<i>"Quick Start: Using the Management Interface Wizard" on page 1-10</i> <i>Chapter 9, "Configuring Secure Socket Layer (SSL)"</i>
SNMP	public, unrestricted	In the default configuration, the switch is open to access by management stations running SNMP management applications capable of viewing and changing the settings and status data in the switch's MIB (Management Information Base). Thus, controlling SNMP access to the switch and preventing unauthorized SNMP access should be a key element of your network security strategy.	<i>"SNMP Security Guidelines" on page 1-15</i> <i>"Quick Start: Using the Management Interface Wizard" on page 1-10</i> <i>Management and Configuration Guide, Chapter 14, refer to the section "Using SNMP Tools To Manage the Switch"</i>
Authorized IP Managers	none	This feature uses IP addresses and masks to determine whether to allow management access to the switch across the network through the following : <ul style="list-style-type: none"> • Telnet and other terminal emulation applications • The switch's Web browser interface • SNMP (with a correct community name) 	<i>Chapter 15, "Using Authorized IP Managers"</i>
Secure Management VLAN	disabled	This feature creates an isolated network for managing the ProCurve switches that offer this feature. When a secure management VLAN is enabled, CLI, Menu interface, and Web browser interface access is restricted to ports configured as members of the VLAN.	<i>Advanced Traffic Management Guide, refer to the chapter "Static Virtual LANs (VLANs)"</i>
TACACS+ Authentication	disabled	This application uses a central server to allow or deny access to TACACS-aware devices in your network. TACACS+ uses username/password sets with associated privilege levels to grant or deny access through either the switch's serial (console) port or remotely, with Telnet. If the switch fails to connect to a TACACS+ server for the necessary authentication service, it defaults to its own locally configured passwords for authentication control. TACACS+ allows both login (read-only) and enable (read/write) privilege level access.	<i>Chapter 5, "TACACS+ Authentication"</i>
RADIUS Authentication	disabled	For each authorized client, RADIUS can be used to authenticate operator or manager access privileges on the switch via the serial port (CLI and Menu interface), Telnet, SSH, and Secure FTP/Secure Copy (SFTP/SCP) access methods.	<i>Chapter 6, "RADIUS Authentication and Accounting"</i>

Feature	Default Setting	Security Guidelines	More Information and Configuration Details
802.1X Access Control	none	<p>This feature provides port-based or user-based authentication through a RADIUS server to protect the switch from unauthorized access and to enable the use of RADIUS-based user profiles to control client access to network services. Included in the general features are the following:</p> <ul style="list-style-type: none">• user-based access control supporting up to 32 authenticated clients per port• port-based access control allowing authentication by a single client to open the port• switch operation as a supplicant for point-to-point connections to other 802.1X-compliant ProCurve switches	<i>Chapter 13 “Configuring Port-Based and User-Based Access Control (802.1X)”</i>
Web and MAC Authentication	none	<p>These options are designed for application on the edge of a network to provide port-based security measures for protecting private networks and the switch itself from unauthorized access. Because neither method requires clients to run any special supplicant software, both are suitable for legacy systems and temporary access situations where introducing supplicant software is not an attractive option.</p> <p>Both methods rely on using a RADIUS server for authentication. This simplifies access security management by allowing you to control access from a master database in a single server. It also means the same credentials can be used for authentication, regardless of which switch or switch port is the current access point into the LAN. Web authentication uses a web page login to authenticate users for access to the network. MAC authentication grants access to a secure network by authenticating device MAC addresses for access to the network.</p>	<i>Chapter 4, “Web and MAC Authentication”</i>

Network Security Features

This section outlines features and defence mechanisms for protecting access through the switch to the network. For more detailed information, see the indicated chapters.

Table 1-2. Network Security—Default Settings and Security Guidelines

Feature	Default Setting	Security Guidelines	More Information and Configuration Details
Secure File Transfers	not applicable	Secure Copy and SFTP provide a secure alternative to TFTP and auto-TFTP for transferring sensitive information such as configuration files and log information between the switch and other devices.	<i>Management and Configuration Guide, Appendix A “File Transfers”, refer to the section “Using Secure Copy and SFTP”</i>
Traffic/Security Filters	none	<p>These statically configured filters enhance in-band security (and improve control over access to network resources) by forwarding or dropping inbound network traffic according to the configured criteria. Filter options include:</p> <ul style="list-style-type: none"> • source-port filters: Inbound traffic from a designated, physical source-port will be forwarded or dropped on a per-port (destination) basis. 	<i>Chapter 12, “Traffic/Security Filters and Monitors”</i>
Port Security, MAC Lockdown, and MAC Lockout	none	<p>The features listed below provide device-based access security in the following ways:</p> <ul style="list-style-type: none"> • Port security: Enables configuration of each switch port with a unique list of the MAC addresses of devices that are authorized to access the network through that port. This enables individual ports to detect, prevent, and log attempts by unauthorized devices to communicate through the switch. Some switch models also include eavesdrop prevention in the port security feature. • MAC lockdown: This “static addressing” feature is used as an alternative to port security to prevent station movement and MAC address “hijacking” by allowing a given MAC address to use only one assigned port on the switch. MAC lockdown also restricts the client device to a specific VLAN. • MAC lockout: This feature enables blocking of a specific MAC address so that the switch drops all traffic to or from the specified address. 	<p><i>Chapter 11, “Configuring and Monitoring Port Security”</i></p> <p>See also “Precedence of Port-Based Security Options” on page 1-17</p>

Feature	Default Setting	Security Guidelines	More Information and Configuration Details
Connection-Rate Filtering based on Virus-Throttling Technology	none	This feature helps protect the network from attack and is recommended for use on the network edge. It is primarily focused on the class of worm-like malicious code that tries to replicate itself by taking advantage of weaknesses in network applications behind unsecured ports. In this case, the malicious code tries to create a large number of outbound connections on an interface in a short time. Connection-Rate filtering detects hosts that are generating traffic that exhibits this behavior, and causes the switch to generate warning messages and (optionally) to throttle or drop all traffic from the offending hosts.	<i>Chapter 3, "Virus Throttling (Connection-Rate Filtering)"</i>
ICMP Rate-Limiting	none	This feature helps defeat ICMP denial-of-service attacks by restricting ICMP traffic to percentage levels that permit necessary ICMP functions, but throttle additional traffic that may be due to worms or viruses (reducing their spread and effect).	<i>Management and Configuration Guide</i> , in the chapter on "Port Traffic Controls" refer to the section "ICMP Rate-Limiting"
Spanning Tree Protection	none	These features prevent your switch from malicious attacks or configuration errors: <ul style="list-style-type: none"> • BPDU Filtering and BPDU Protection: Protects the network from denial-of-service attacks that use spoofing BPDUs by dropping incoming BPDUs frames and/or blocking traffic through a port. • STP Root Guard: Protects the STP root bridge from malicious attacks or configuration mistakes. 	<i>Advanced Traffic Management Guide</i> , refer to the chapter "Multiple Instance Spanning-Tree Operation"
DHCP Snooping, Dynamic ARP Protection, and Dynamic IP Lockdown	none	These features provide the following additional protections for your network: <ul style="list-style-type: none"> • DHCP Snooping: Protects your network from common DHCP attacks, such as address spoofing and repeated address requests. • Dynamic ARP Protection: Protects your network from ARP cache poisoning. • Dynamic IP Lockdown: Prevents IP source address spoofing on a per-port and per-VLAN basis • Instrumentation Monitor. Helps identify a variety of malicious attacks by generating alerts for detected anomalies on the switch. 	<i>Chapter 11, "Configuring Advanced Threat Protection"</i>

Getting Started with Access Security

ProCurve switches are designed as “plug and play” devices, allowing quick and easy installation in your network. In its default configuration the switch is open to unauthorized access of various types. When preparing the switch for network operation, therefore, ProCurve strongly recommends that you enforce a security policy to help ensure that the ease in getting started is not used by unauthorized persons as an opportunity for access and possible malicious actions.

Since security incidents can originate with sources inside as well as outside of an organization, your access security provisions must protect against internal and external threats while preserving the necessary network access for authorized clients and users. It is important to evaluate the level of management access vulnerability existing in your network and take steps to ensure that all reasonable security precautions are in place. This includes both configurable security options and physical access to the switch.

Switch management access is available through the following methods:

- Front panel access to the console serial port (see “Physical Security”)
- Inbound Telnet access
- Web-browser access
- SNMP access

For guidelines on locking down your switch for remote management access, see “Quick Start: Using the Management Interface Wizard” on page 1-10.

Physical Security

Physical access to the switch allows the following:

- use of the console serial port (CLI and Menu interface) for viewing and changing the current configuration and for reading status, statistics, and log messages.
- use of the switch’s USB port for file transfers and autorun capabilities.
- use of the switch’s Clear and Reset buttons for these actions:
 - clearing (removing) local password protection
 - rebooting the switch
 - restoring the switch to the factory default configuration (and erasing any non-default configuration settings)

Keeping the switch in a locked wiring closet or other secure space helps to prevent unauthorized physical access.

As additional precautions, you can do the following:

- Disable or re-enable the password-clearing function of the Clear button.
- Configure the Clear button to reboot the switch after clearing any local usernames and passwords.
- Modify the operation of the Reset+Clear button combination so that the switch reboots, but does not restore the switch's factory default settings.
- Disable or re-enable password recovery.

For the commands used to configure the Clear and Reset buttons, refer to “Front-Panel Security” on page 2-23.

Quick Start: Using the Management Interface Wizard

The Management Interface wizard provides a convenient step-by-step method to prepare the switch for secure network operation. It guides you through the process of locking down the following switch operations or protocols:

- setting local passwords
- restricting SNMP access
- enabling/disabling Telnet
- enabling/disabling SSH
- enabling/disabling remote Web management
- restricting web access to SSL
- setting timeouts for SSH/Telnet sessions

The wizard can also be used to view the pre-configured defaults and see the current settings for switch access security. The wizard can be launched either via the CLI (see page 1-10) or the Web browser interface (see page 1-12).

Note

The wizard's security settings can also be configured using standard commands via the CLI, Menu or Web browser interfaces. For full details on preparing and configuring the switch for SSH and SSL operation, refer to chapters 6 and 7 respectively.

CLI: Management Interface Wizard

To configure security settings using the CLI wizard, follow the steps below:

1. At the command prompt, type **setup mgmt-interfaces**.

The welcome banner appears and the first setup option is displayed (**Operator password**). As you advance through the wizard, each setup option displays the current value in brackets [] as shown in Figure 1-1.

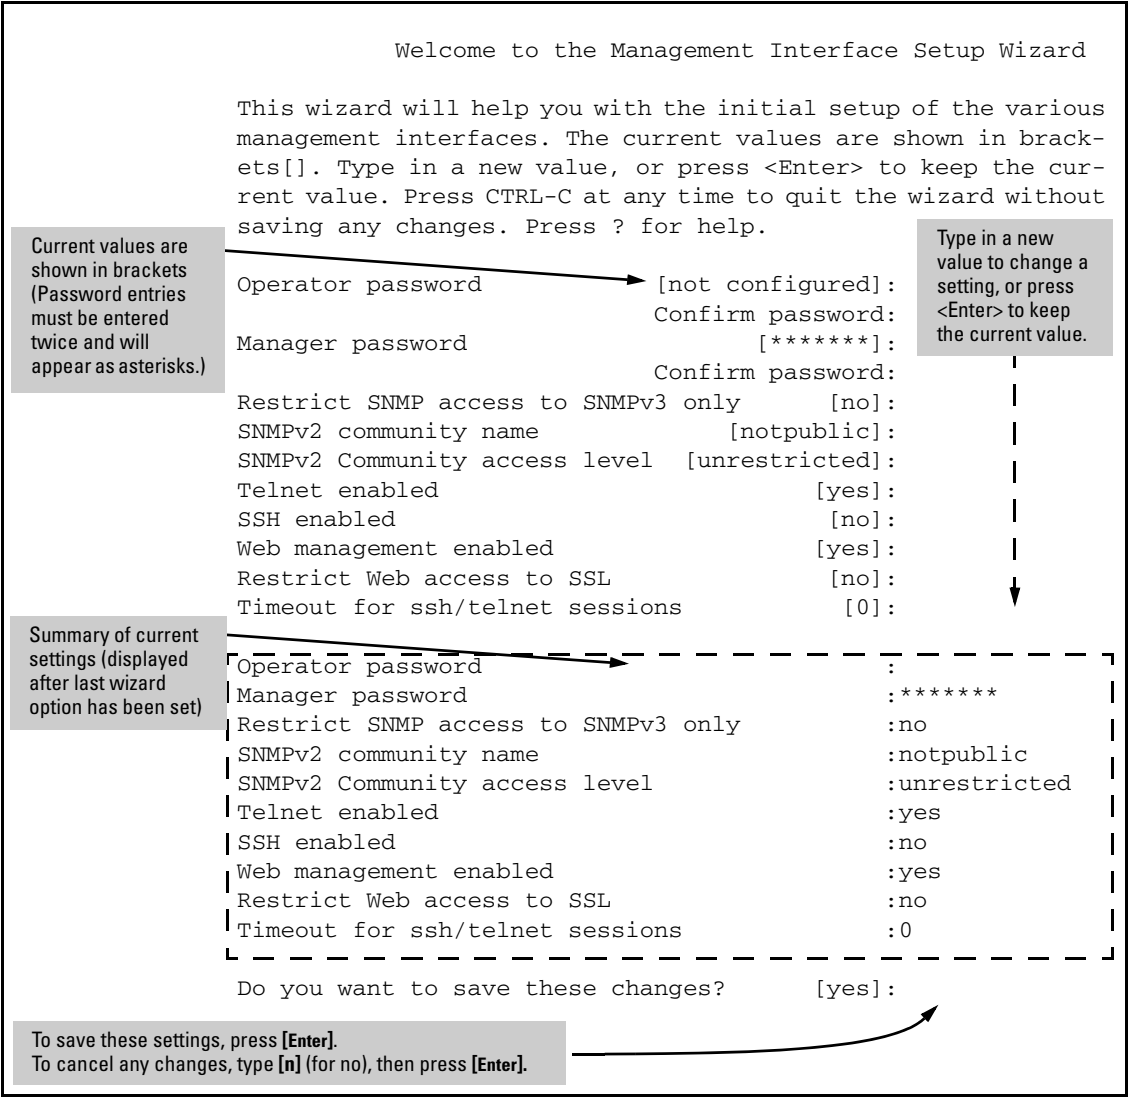


Figure 1-1. Example of Management Interface Wizard Configuration

2. When you enter the wizard, you have the following options:
 - To update a setting, type in a new value, or press **[Enter]** to keep the current value.
 - To quit the wizard without saving any changes, press **[CTRL-C]** at any time.
 - To access online Help for any option, press **[?]**.

After you have gone through each setup option, the wizard displays the summary configuration together with a prompt to save the changes (see Figure 1-1 on page 1-11 for an example).

3. When the message appears asking if you want to save these changes, you have the following options:
 - To save your changes, press **[Enter]**.
 - To cancel any changes without saving, type **[n]** and then press **[Enter]**.

After pressing **[Enter]**, the wizard exits to the command line prompt.

CLI Wizard: Operating Notes and Restrictions.

- Once a password has been configured on the switch, you cannot remove it using the CLI wizard. Passwords can be removed by executing the **no password** command directly from the CLI.
- When you restrict SNMP access to SNMPv3 only, the options SNMPv2 community name and access level will not appear.
- The wizard displays the first available SNMPv2 community and allows the user to modify the first community access parameters.
- The wizard creates a new SNMP community only when no communities have been configured on the switch.

Web: Management Interface Wizard

To use the Management Interface wizard from the Web browser interface, follow the steps below:

1. Click the **Security** tab.
2. Click the **Setup Wizard** button.

The Welcome window appears.

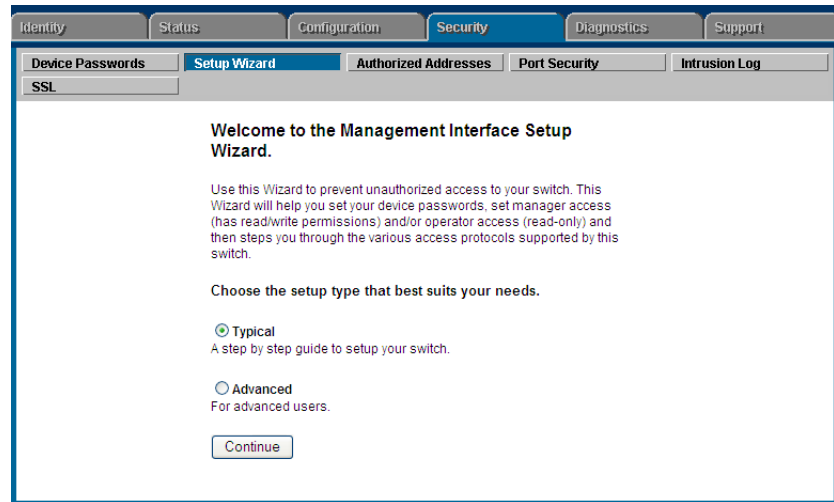


Figure 1-2. Management Interface Wizard: Welcome Window

This page allows you to choose between two setup types:

- **Typical**—provides a multiple page, step-by-step method to configure security settings, with on-screen instructions for each option.
 - **Advanced**—provides a single summary screen in which to configure all security settings at once.
3. To enter the wizard, choose a setup option and then click **Continue**.
- **Typical** (multi-page setup): when you select this option, you will get an alert indicating that configuration changes will be lost if you click on the Web browser's navigation tabs. Click **OK** to close the alert and then advance through the following setup pages: **Operator Password**, **Manager Password**, **SNMP**, **Telnet**, **SSH**, **Web Management GUI**, **Timeout** (see pages for details and setup options).

At each page, you have the following options:

- Enter a new value and/or make a new selection, or click **Continue** to keep the current value and proceed to the next page setup. When you have gone through all configurable options, the summary setup page appears (see step 4).
 - To quit the Setup and return to the Welcome window without saving any changes, click **Exit** at any time.
 - To return to the previous screen(s), click **Back** at any time.
- **Advanced** (single page setup): when you select this option, the summary setup page appears immediately (see step 4).

4. The summary setup screen displays the current configuration settings for all setup options (see Figure 1-3).

The screenshot shows a web interface for a network device. At the top, there are tabs: Identity, Status, Configuration, Security (selected), Diagnostics, and Support. Below these are sub-tabs: Device Passwords, Setup Wizard (selected), Authorized Addresses, Port Security, and Intrusion Log. The main content area is titled "Management Interface Setup." and contains three sections:

- Passwords:**
 - Operator password: [text input]
 - Confirm Password: [text input]
 - Manager password: [password input]
 - Confirm Password: [password input]
- SNMP Access:**
 - ☐ Restrict SNMP access to SNMPv3 only
 - SNMPv2 Community Name: [text input with value "public"]
 - SNMPv2 Community Access Level: ☐ Restricted ☒ Unrestricted
 - SNMPv2 Operator/Manager Access: ☐ Operator ☒ Manager
- Other Access:**
 - ☒ Enable Telnet
 - ☒ Enable SSH
 - ☒ Enable Web
 - ☐ Restrict Web Access to SSL (No Certificate Installed. To install a certificate choose the SSL tab.)
 - ☐ Enable USB Autorun
 - Session timeout: [text input with value "0"]

At the bottom, there are three buttons: Back, Apply, and Exit.

Figure 1-3. Management Interface Wizard: Summary Setup

From this screen, you have the following options:

- To change any setting that is shown, type in a new value or make a different selection.
- To apply the settings permanently, click **Apply**.
- To quit the Setup screen without saving any changes, click **Exit**.
- To return to the previous screen, click **Back**.

Web Wizard: Operating Notes and Restrictions.

- If you click on the Web interface's navigation tab during setup, all configuration changes will be lost.
- When you restrict SNMP access to SNMPv3 only, the SNMPv2 options are not available.
- The option to restrict Web Access to SSL is made available only if a server certificate has been previously installed on the switch.

SNMP Security Guidelines

In the default configuration, the switch is open to access by management stations running SNMP (Simple Network Management Protocol) management applications capable of viewing and changing the settings and status data in the switch's MIB (Management Information Base). Thus, controlling SNMP access to the switch and preventing unauthorized SNMP access should be a key element of your network security strategy.

General SNMP Access to the Switch. The switch supports SNMP versions 1, 2c, and 3, including SNMP community and trap configuration. The default configuration supports versions 1 and 2c compatibility, which uses plain text and does not provide security options.

ProCurve recommends that you enable SNMP version 3 for improved security. SNMPv3 includes the ability to configure restricted access and to block all non-version 3 messages (which blocks version 1 and 2c unprotected operation).

SNMPv3 security options include:

- configuring device communities as a means for excluding management access by unauthorized stations
- configuring for access authentication and privacy
- reporting events to the switch CLI and to SNMP trap receivers
- restricting non-SNMPv3 agents to either read-only access or no access
- co-existing with SNMPv1 and v2c if necessary

SNMP Access to the Authentication Configuration MIB. A management station running an SNMP networked device management application, such as ProCurve Manager Plus (PCM+) or HP OpenView, can access the switch's management information base (MIB) for read access to the switch's status and read/write access to the switch's authentication configuration (hpSwitchAuth). This means that the switch's default configuration now allows SNMP access to security settings in hpSwitchAuth.

**Note on SNMP
Access to
Authentication
MIB**

Downloading and booting from the software for the first time enables SNMP access to the authentication configuration MIB (the default action). If SNMPv3 and other security safeguards are not in place, the switch's authentication configuration MIB is exposed to unprotected SNMP access and you should use the command shown below to disable this access.

If SNMP access to the hpSwitchAuth MIB is considered a security risk in your network, then you should implement the following security precautions when downloading and booting from the software:

- If SNMP access to the authentication configuration (hpSwitchAuth) MIB described above is not desirable for your network, then immediately after downloading and booting from the software for the first time, use the following command to disable this feature:

snmp-server mib hpswitchauthmib excluded

- If you choose to leave the authentication configuration MIB accessible, then you should do the following to help ensure that unauthorized workstations cannot use SNMP tools to access the MIB:
 - a. Configure SNMP version 3 management and access security on the switch.
 - b. Disable SNMP version 2c on the switch.

For details on this feature, refer to the section titled “Using SNMP To View and Configure Switch Authentication Features” on page 5-21.

For more information on configuring SNMP, refer to the section “*Using SNMP Tools To Manage the Switch*” in the chapter “*Configuring for Network Management Applications*” in the *Management and Configuration Guide* for your switch.

Precedence of Security Options

This section explains how port-based security options, and client-based attributes used for authentication, get prioritized on the switch.

Precedence of Port-Based Security Options

Where the switch is running multiple security options, it implements network traffic security based on the OSI (Open Systems Interconnection model) precedence of the individual options, from the lowest to the highest. The following list shows the order in which the switch implements configured security features on traffic moving through a given port.

1. Disabled/Enabled physical port
 2. MAC lockout (Applies to all ports on the switch.)
 3. MAC lockdown
 4. Port security
 5. Authorized IP Managers
 6. Application features at higher levels in the OSI model, such as SSH.
- (The above list does not address the mutually exclusive relationship that exists among some security features.)

Precedence of Client-Based Authentication: Dynamic Configuration Arbiter

The Dynamic Configuration Arbiter (DCA) is implemented to determine the client-specific parameters that are assigned in an authentication session.

A client-specific authentication configuration is bound to the MAC address of a client device and may include the following parameters:

- Untagged client VLAN ID
- Tagged VLAN IDs
- Per-port CoS (802.1p) priority
- Per-port rate-limiting on inbound traffic

DCA allows client-specific parameters configured in any of the following ways to be applied and removed as needed in a specified hierarchy of precedence. When multiple values for an individual configuration parameter exist, the

value applied to a client session is determined in the following order (from highest to lowest priority) in which a value configured with a higher priority overrides a value configured with a lower priority:

1. Attribute profiles applied through the Network Immunity network-management application using SNMP (see “Network Immunity Manager”)
2. 802.1X authentication parameters (RADIUS-assigned)
3. Web- or MAC-authentication parameters (RADIUS-assigned)
4. Local, statically-configured parameters

Although RADIUS-assigned settings are never applied to ports for non-authenticated clients, the Dynamic Configuration Arbiter allows you to configure and assign client-specific port configurations to non-authenticated clients, provided that a client’s MAC address is known in the switch in the forwarding database. DCA arbitrates the assignment of attributes on both authenticated and non-authenticated ports.

DCA does not support the arbitration and assignment of client-specific attributes on trunk ports.

Network Immunity Manager

Network Immunity Manager (NIM) is a plug-in to ProCurve Manager (PCM) and a key component of the ProCurve Network Immunity security solution that provides comprehensive detection and per-port-response to malicious traffic at the ProCurve network edge. NIM allows you to apply policy-based actions to minimize the negative impact of a client’s behavior on the network. For example, using NIM you can apply a client-specific profile that adds or modifies per-port rate-limiting and VLAN ID assignments.

Note

NIM actions only support the configuration of per-port rate-limiting and VLAN ID assignment; NIM does not support CoS (802.1p) priority assignment.

NIM-applied parameters temporarily override RADIUS-configured and locally configured parameters in an authentication session. When the NIM-applied action is removed, the previously applied client-specific parameter (locally configured or RADIUS-assigned) is re-applied unless there have been other configuration changes to the parameter. In this way, NIM allows you to minimize network problems without manual intervention.

NIM also allows you to configure and apply client-specific profiles on ports that are not configured to authenticate clients (unauthorized clients), provided that a client’s MAC address is known in the switch’s forwarding database.

The profile of attributes applied for each client (MAC address) session is stored in the `hpicfUsrProfile` MIB, which serves as the configuration interface for Network Immunity Manager. A client profile consists of NIM-configured, RADIUS-assigned, and statically configured parameters. Using **show** commands for 802.1X, web or MAC authentication, you can verify which RADIUS -assigned and statically configured parameters are supported and if they are supported on a per-port or per-client basis.

A NIM policy accesses the `hpicfUsrProfileMIB` through SNMP to perform the following actions:

- Bind (or unbind) a profile of configured attributes to the MAC address of a client device on an authenticated or unauthenticated port.
- Configure or unconfigure an untagged VLAN for use in an authenticated or unauthenticated client session.

Note that the attribute profile assigned to a client is often a combination of NIM-configured, RADIUS-assigned, and statically configured settings. Precedence is always given to the temporarily applied NIM-configured parameters over RADIUS-assigned and locally configured parameters.

For information on Network Immunity Manager, go to the HP ProCurve Networking Web site at www.procurve.com/solutions, click on **Security**, and then click on **Security Products**.

Arbitrating Client-Specific Attributes

In previous releases, client-specific authentication parameters for 802.1X Web, and MAC authentication are assigned to a port using different criteria. A RADIUS-assigned parameter is always given highest priority and overrides statically configured local passwords. 802.1X authentication parameters override Web or MAC authentication parameters.

DCA stores three levels of client-specific authentication parameters and prioritizes them according to the following hierarchy of precedence:

1. NIM access policy (applied through SNMP)
2. RADIUS-assigned
 - a. 802.1X authentication
 - b. Web or MAC authentication
3. Statically (local) configured

Client-specific configurations are applied on a per-parameter basis on a port. In a client-specific profile, if DCA detects that a parameter has configured values from two or more levels in the hierarchy of precedence described above, DCA decides which parameters to add or remove, or whether to fail the authentication attempt due to an inability to apply the parameters.

For example, NIM may configure only rate-limiting for a specified client session, while RADIUS-assigned values may include both an untagged VLAN ID and a rate-limiting value to be applied. In this case, DCA applies the NIM-configured rate-limiting value and the RADIUS-assigned VLAN (if there are no other conflicts).

Also, you can assign NIM-configured parameters (for example, VLAN ID assignment or rate-limiting) to be activated in a client session when a threat to network security is detected. When the NIM-configured parameters are later removed, the parameter values in the client session return to the RADIUS-configured or locally configured settings, depending on which are next in the hierarchy of precedence.

In addition, DCA supports conflict resolution for QoS (port-based CoS priority) and rate-limiting (ingress) by determining whether to configure either strict or non-strict resolution on a switch-wide basis. For example, if multiple clients authenticate on a port and a rate-limiting assignment by a newly authenticating client conflicts with the rate-limiting values assigned to previous clients, by using Network Immunity you can configure the switch to apply any of the following attributes:

- Apply only the latest rate-limiting value assigned to all clients.
- Apply a client-specific rate-limiting configuration to the appropriate client session (overwrites any rate-limit previously configured for other client sessions on the port).

For information about how to configure RADIUS-assigned and locally configured authentication settings, refer to:

- RADIUS-assigned 802.1X authentication: “*Configuring Port-Based and User-Based Access Control (802.1X)*” on page 10-1.
- RADIUS-assigned Web or MAC authentication: “*Web and MAC Authentication*” on page 3-1.
- RADIUS-assigned CoS, and rate-limiting: “*Configuring RADIUS Server Support for Switch Services*” on page 7-1.
- Statically (local) configured: “*Configuring Username and Password Security*” on page 2-1.

ProCurve Identity-Driven Manager (IDM)

IDM is a plug-in to ProCurve Manager Plus (PCM+) and uses RADIUS-based technologies to create a user-centric approach to network access management and network activity tracking and monitoring. IDM enables control of access security policy from a central management server, with policy enforcement to the network edge, and protection against both external and internal threats.

Using IDM, a system administrator can configure automatic and dynamic security to operate at the network edge when a user connects to the network. This operation enables the network to:

- approve or deny access at the edge of the network instead of in the core;
- distinguish among different users and what each is authorized to do;
- configure guest access without compromising internal security.

Criteria for enforcing RADIUS-based security for IDM applications includes classifiers such as:

- authorized user identity
- authorized device identity (MAC address)
- software running on the device
- physical location in the network
- time of day

Responses can be configured to support the networking requirements, user (SNMP) community, service needs, and access security level for a given client and device.

For more information on IDM, go to the ProCurve Web site at www.procurve.com/solutions, click on **Security**, and then click **Security Products**.

Configuring Username and Password Security

Contents

Overview	2-3
Configuring Local Password Security	2-6
Menu: Setting Passwords	2-6
CLI: Setting Passwords and Usernames	2-8
Web: Setting Passwords and Usernames	2-9
SNMP: Setting Passwords and Usernames	2-9
Saving Security Credentials in a Config File	2-10
Benefits of Saving Security Credentials	2-10
Enabling the Storage and Display of Security Credentials	2-11
Security Settings that Can Be Saved	2-11
Local Manager and Operator Passwords	2-12
Password Command Options	2-12
SNMP Security Credentials	2-13
802.1X Port-Access Credentials	2-14
TACACS+ Encryption Key Authentication	2-15
RADIUS Shared-Secret Key Authentication	2-15
SSH Client Public-Key Authentication	2-16
Operating Notes	2-19
Restrictions	2-21
Front-Panel Security	2-23
When Security Is Important	2-23
Front-Panel Button Functions	2-24
Clear Button	2-25
Reset Button	2-25
Restoring the Factory Default Configuration	2-25
Configuring Front-Panel Security	2-27
Disabling the Clear Password Function of the Clear Button	2-29

Re-Enabling the Clear Button and Setting or	
Changing the “Reset-On-Clear” Operation	2-30
Changing the Operation of the Reset+Clear Combination	2-31
Password Recovery	2-32
Disabling or Re-Enabling the Password Recovery Process	2-32
Password Recovery Process	2-34

Overview

Feature	Default	Menu	CLI	Web
Set Usernames	none	—	—	page 2-9
Set a Password	none	page 2-6	page 2-8	page 2-9
Delete Password Protection	n/a	page 2-7	page 2-8	page 2-9
show front-panel-security	n/a	—	page 1-13	—
front-panel-security		—	page 1-13	—
password-clear	enabled	—	page 1-13	—
reset-on-clear	disabled	—	page 1-14	—
factory-reset	enabled	—	page 1-15	—
password-recovery	enabled	—	page 1-15	—

Console access includes both the menu interface and the CLI. There are two levels of console access: Manager and Operator. For security, you can set a *password pair* (username and password) on each of these levels.

Notes

Username are optional. Also, in the menu interface, you can configure passwords, but not usernames. To configure usernames, use the CLI or the web browser interface.

Username and passwords for Manager and Operator access can also be configured using SNMP. For more information, refer to “Using SNMP To View and Configure Switch Authentication Features” on page 5-21.

Username and passwords for Manager and Operator access can also be configured using the Management Interface Wizard. For more information, refer to “Quick Start: Using the Management Interface Wizard” on page 1-10.

Level	Actions Permitted
Manager:	Access to all console interface areas. <i>This is the default level.</i> That is, if a Manager password has <i>not</i> been set prior to starting the current console session, then anyone having access to the console can access any area of the console interface.
Operator:	Access to the Status and Counters menu, the Event Log, and the CLI*, but no Configuration capabilities. On the Operator level, the configuration menus, Download OS, and Reboot Switch options in the Main Menu are not available.
*Allows use of the ping, link-test, show, menu, exit, and logout commands, plus the enable command if you can provide the Manager password.	

To configure password security:

1. Set a Manager password pair (and an Operator password pair, if applicable for your system).
2. Exit from the current console session. A Manager password pair will now be needed for full access to the console.

If you do steps 1 and 2, above, then the next time a console session is started for either the menu interface or the CLI, a prompt appears for a password. Assuming you have protected both the Manager and Operator levels, the level of access to the console interface will be determined by which password is entered in response to the prompt.

If you set a Manager password, you may also want to configure an inactivity timer. This causes the console session to end after the specified period of inactivity, thus giving you added security against unauthorized console access. You can use either of the following to set the inactivity timer:

- **Menu Interface:** System Information screen (Select “2. Switch Configuration.”)
- **CLI:** Use the **console inactivity-timer** **< 0 | 1 | 5 | 10 | 15 | 20 | 30 | 60 | 120 >**

Notes

The manager and operator passwords and (optional) usernames control access to the menu interface, CLI, and web browser interface.

If you configure only a Manager password (with no Operator password), and in a later session the Manager password is not entered correctly in response to a prompt from the switch, then the switch does not allow management access for that session.

If the switch has a password for both the Manager and Operator levels, and neither is entered correctly in response to the switch's password prompt, then the switch does not allow management access for that session.

Passwords are case-sensitive.

When configuring an operator or manager password a message will appear indicating that (USB) autorun has been disabled. For more information on the autorun feature, refer to the Appendix A on "File Transfers" in the *Management and Configuration Guide* for your switch.

Caution

If the switch has neither a Manager nor an Operator password, anyone having access to the switch through either Telnet, the serial port, or the web browser interface can access the switch with full manager privileges. Also, if you configure only an Operator password, entering the Operator password enables full manager privileges.

The rest of this chapter covers how to:

- Set passwords
- Delete passwords
- Recover from a lost password
- Maintain front-panel security

Configuring Local Password Security

Menu: Setting Passwords

As noted earlier in this section, usernames are optional. Configuring a username requires either the CLI or the web browser interface.

1. From the Main Menu select:

3. Console Passwords

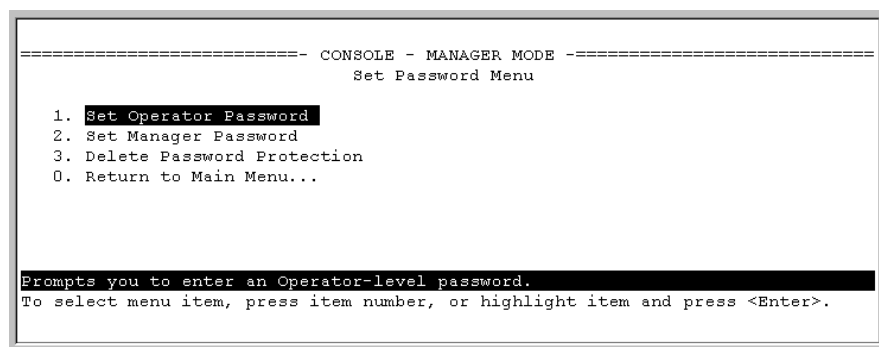


Figure 2-1. The Set Password Screen

2. To set a new password:
 - a. Select **Set Manager Password** or **Set Operator Password**. You will then be prompted with **Enter new password**.
 - b. Type a password of up to 16 ASCII characters with no spaces and press **[Enter]**. (Remember that passwords are case-sensitive.)
 - c. When prompted with **Enter new password again**, retype the new password and press **[Enter]**.

After you configure a password, if you subsequently start a new console session, you will be prompted to enter the password. (If you use the CLI or web browser interface to configure an optional username, the switch will prompt you for the username, and then the password.)

To Delete Password Protection (Including Recovery from a Lost Password): This procedure deletes *all* usernames (if configured) and passwords (Manager and Operator).

If you have physical access to the switch, press and hold the Clear button (on the front of the switch) for a minimum of one second to clear all password protection, then enter new passwords as described earlier in this chapter.

If you do not have physical access to the switch, you will need Manager-Level access:

1. Enter the console at the Manager level.
2. Go to the **Set Passwords** screen as described above.
3. Select **Delete Password Protection**. You will then see the following prompt:

Continue Deletion of password protection? No

4. Press the Space bar to select **Yes**, then press **[Enter]**.
5. Press **[Enter]** to clear the Password Protection message.

To Recover from a Lost Manager Password: If you cannot start a console session at the Manager level because of a lost Manager password, you can clear the password by getting physical access to the switch and pressing and holding the Clear button for a minimum of one second. This action deletes all passwords and usernames (Manager and Operator) used by both the console and the web browser interface.

CLI: Setting Passwords and Usernames

Commands Used in This Section

password	See below.
----------	------------

Configuring Manager and Operator Passwords.

Note

You can configure manager and operator passwords in one step. See “Saving Security Credentials in a Config File” on page 2-10 of this guide.

Syntax: [no] password <manager | operator | all | port-access>
[user-name ASCII-STR] [<plaintext | sha1> ASCII-STR]

```
ProCurve(config)# password manager
New password: *
Please retype new password: *
ProCurve(config)# password operator
New password: *
Please retype new password: *
```

• Password entries appear as asterisks.

• You must type the password entry twice.

Figure 2-2. Example of Configuring Manager and Operator Passwords

To Remove Password Protection. Removing password protection means to eliminate password security. This command prompts you to verify that you want to remove one or both passwords, then clears the indicated password(s). (This command also clears the username associated with a password you are removing.) For example, to remove the Operator password (and username, if assigned) from the switch, you would do the following:

```
ProCurve(config)# no password
Password protection will be deleted, do you want to continue [y/n]? y
ProCurve(config)#
```

Press [Y] (for yes) and press [Enter].

Figure 2-3. Removing a Password and Associated Username from the Switch

The effect of executing the command in figure 2-3 is to remove password protection from the Operator level. (This means that anyone who can access the switch console can gain Operator access without having to enter a username or password.)

If you want to remove both operator and manager password protection, use the **no password all** command.

Web: Setting Passwords and Usernames

In the web browser interface you can enter passwords and (optional) usernames.

To Configure (or Remove) Usernames and Passwords in the Web Browser Interface.

1. Click on the **Security** tab.

Click on **[Device Passwords]**.

2. Do one of the following:
 - To set username and password protection, enter the usernames and passwords you want in the appropriate fields.
 - To remove username and password protection, leave the fields blank.
3. Implement the usernames and passwords by clicking on **[Apply Changes]**.

SNMP: Setting Passwords and Usernames

Usernames and passwords for Manager and Operator access can also be configured using SNMP. For more information, refer to “Using SNMP To View and Configure Switch Authentication Features” on page 5-21.

Saving Security Credentials in a Config File

You can store and view the following security settings in internal flash memory by entering the **include-credentials** command:

- Local manager and operator passwords and (optional) user names that control access to a management session on the switch through the CLI, menu interface, or web browser interface
- SNMP security credentials used by network management stations to access a switch, including authentication and privacy passwords
- Port-access passwords and usernames used as 802.1X authentication credentials for access to the switch
- TACACS+ encryption keys used to encrypt packets and secure authentication sessions with TACACS+ servers
- RADIUS shared secret (encryption) keys used to encrypt packets and secure authentication sessions with RADIUS servers
- Secure Shell (SSH) public keys used to authenticate SSH clients that try to connect to the switch.

Benefits of Saving Security Credentials

The benefits of including and saving security credentials are as follows:

- After making changes to security parameters in the running configuration, you can experiment with the new configuration and, if necessary, view the new security settings during the session. After verifying the configuration, you can then save it permanently by writing the settings to the startup-config file.
- By permanently saving a switch's security credentials in internal flash memory, you can upload the file to a TFTP server, and later download the file to the ProCurve switches on which you want to use the same security settings without having to manually configure the settings (except for SNMPv3 user parameters) on each switch.
- By storing different security settings in different files, you can test different security configurations when you first download a new software version that supports multiple configuration files, by changing the configuration file used when you reboot the switch.

For more information about how to experiment with, upload, download, and use configuration files with different software versions, refer to the following:

- The chapter on “Switch Memory and Configuration” in the *Management and Configuration Guide*.
- “Configuring Local Password Security” on page 2-6 in this guide.

Enabling the Storage and Display of Security Credentials

To enable the security settings, enter the **include-credentials** command.

Syntax: [no] include-credentials

Enables the inclusion and display of the currently configured manager and operator usernames and passwords, RADIUS shared secret keys, SNMP and 802.1X authenticator (port-access) security credentials, and SSH client public-keys in the running configuration.

To view the currently configured security settings in the running configuration, enter one of the following commands:

- **show running-config:** *Displays the configuration settings in the current running-config file.*
- **write terminal:** *Displays the configuration settings in the current running-config file.*

For more information, refer to “Switch Memory and Configuration” in the Management and Configuration Guide.

The “no” form of the command disables only the display and copying of these security parameters from the running configuration, while the security settings remain active in the running configuration.

Default: The security credentials described in “Security Settings that Can Be Saved” on page 2-11 are not stored in the running configuration.

Security Settings that Can Be Saved

The security settings that can be saved are:

- Local manager and operator passwords and user names
- SNMP security credentials, including SNMPv1 community names and SNMPv3 usernames, authentication, and privacy settings
- 802.1X port-access passwords and usernames
- TACACS+ encryption keys
- RADIUS shared secret (encryption) keys
- Public keys of SSH-enabled management stations that are used by the switch to authenticate SSH clients that try to connect to the switch

Local Manager and Operator Passwords

The information saved to the running-config file when the **include-credentials** command is entered includes:

```
password manager [user-name <name>] <hash-type> <pass-hash>  
password operator [user-name <name>] <hash-type> <pass-hash>
```

where

<name> is an alphanumeric string for the user name assigned to the manager or operator.

<hash-type> indicates the type of hash algorithm used: SHA-1 or plain text.

<pass-hash> is the SHA-1 authentication protocol's hash of the password or clear ASCII text.

For example, a manager username and password may be stored in a running-config file as follows:

```
password manager user-name George SHA1  
2fd4e1c67a2d28fced849ee1bb76e7391b93eb12
```

Use the **write memory** command to save the password configurations in the startup-config file. The passwords take effect when the switch boots with the software version associated with that configuration.

Caution

If a startup configuration file includes other security credentials, but does not contain a manager or operator password, the switch will not have password protection and can be accessed through Telnet, the serial port, or web interface with full manager privileges.

Password Command Options

The **password** command has the following options:

Syntax: [no] password <manager | operator | port-access> [user-name <name>]
 <hash-type> <password>

Set or clear a local username/password for a given access level.

manager: *configures access to the switch with manager-level privileges.*

operator: *configures access to the switch with operator-level privileges.*

port-access: *configures access to the switch through 802.1X authentication with operator-level privileges.*

user-name <name>: the optional text string of the user name associated with the password.

<hash-type>: specifies the type of algorithm (if any) used to hash the password. Valid values are **plaintext** or **sha-1**

<password>: the clear ASCII text string or SHA-1 hash of the password.

You can enter a manager, operator, or 802.1X port-access password in clear ASCII text or hashed format. However, manager and operator passwords are displayed and saved in a configuration file only in hashed format; port-access passwords are displayed and saved only as plain ASCII text.

After you enter the complete command syntax, the password is set. You are not prompted to enter the password a second time.

This command enhancement allows you to configure manager, operator, and 802.1X port-access passwords in only one step (instead of entering the **password** command and then being prompted twice to enter the actual password).

- For more information about configuring local manager and operator passwords, refer to “Configuring Username and Password Security” on page 2-1 in this guide.
- For more information about configuring a port-access password for 802.1X client authentication, see “802.1X Port-Access Credentials” on page 2-14 in this guide.

SNMP Security Credentials

SNMPv1 community names and write-access settings, and SNMPv3 usernames continue to be saved in the running configuration file even when you enter the **include-credentials** command.

In addition, the following SNMPv3 security parameters are also saved:

```
snmpv3 user "<name>" [auth <md5|sha> "<auth-pass>"]  
[priv "<priv-pass>"]
```

where:

<name> is the name of an SNMPv3 management station.

[auth <md5|sha>] is the (optional) authentication method used for the management station.

<auth-pass> is the hashed authentication password used with the configured authentication method.

[**priv** <*priv-pass*>] is the (optional) hashed privacy password used by a privacy protocol to encrypt SNMPv3 messages between the switch and the station.

The following example shows the additional security credentials for SNMPv3 users that can be saved in a running-config file:

```
snmpv3 user boris \  
auth md5 "9e4cfef901f21cf9d21079debeca453" \  
priv "82ca4dc99e782db1a1e914f5d8f16824"  
  
snmpv3 user alan \  
auth sha "8db06202b8f293e9bc0c00ac98cf91099708ecdf" \  
priv "5bc4313e9fd7c2953aaea9406764fe8bb629a538"
```

Figure 2-4. Example of Security Credentials Saved in the Running-Config

Although you can enter an SNMPv3 authentication or privacy password in either clear ASCII text or the SHA-1 hash of the password, the password is displayed and saved in a configuration file only in hashed format, as shown in the preceding example.

For more information about the configuration of SNMP security parameters, refer to the chapter on “Configuring for Network Management Applications” in the *Management and Configuration Guide* for your switch.

802.1X Port-Access Credentials

802.1X authenticator (port-access) credentials can be stored in a configuration file. 802.1X *authenticator* credentials are used by a port to authenticate supplicants requesting a point-to-point connection to the switch. 802.1X *supplicant* credentials are used by the switch to establish a point-to-point connection to a port on another 802.1X-aware switch. Only 802.1X authenticator credentials are stored in a configuration file. For information about how to use 802.1X on the switch both as an authenticator and a supplicant, see “Configuring Port-Based and Client-Based Access Control (802.1X)” in this guide.

The local password configured with the **password** command is no longer accepted as an 802.1X authenticator credential. A new configuration command (**password port-access**) is introduced to configure the local operator username and password used as 802.1X authentication credentials for access to the switch.

The **password port-access** values are configured separately from the manager and operator passwords configured with the **password manager** and **password operator** commands and used for management access to the switch. For information on the new **password** command syntax, see “Password Command Options” on page 2-12.

After you enter the complete **password port-access** command syntax, the password is set. You are not prompted to enter the password a second time.

TACACS+ Encryption Key Authentication

You can use TACACS+ servers to authenticate users who request access to a switch through Telnet (remote) or console (local) sessions. TACACS+ uses an authentication hierarchy consisting of:

- Remote passwords assigned in a TACACS+ server
- Local manager and operator passwords configured on the switch.

When you configure TACACS+, the switch first tries to contact a designated TACACS+ server for authentication services. If the switch fails to connect to any TACACS+ server, it defaults to its own locally assigned passwords for authentication control if it has been configured to do so.

For improved security, you can configure a global or server-specific encryption key that encrypts data in TACACS+ packets transmitted between a switch and a RADIUS server during authentication sessions. The key configured on the switch must match the encryption key configured in each

TACACS+ server application. (The encryption key is sometimes referred to as “shared secret” or “secret” key.) For more information, see “TACACS+ Authentication” on page 5-1 in this guide.

TACACS+ shared secret (encryption) keys can be saved by entering this command:

```
ProCurve(config)# tacacs-server key <keystring>
```

The option *<keystring>* is the encryption key (in clear text) used for secure communication with all or a specific TACACS+ server.

RADIUS Shared-Secret Key Authentication

You can use RADIUS servers as the primary authentication method for users who request access to a switch through Telnet, SSH, Web interface, console, or port-access (802.1X). The shared secret key is a text string used to encrypt data in RADIUS packets transmitted between a switch and a RADIUS server

during authentication sessions. Both the switch and the server have a copy of the key; the key is never transmitted across the network. For more information, refer to “3. Configure the Switch To Access a RADIUS Server” on page 6-14 in this guide.

RADIUS shared secret (encryption) keys can be saved in a configuration file by entering this command:

```
ProCurve(config)# radius-server key <keystring>
```

The option *<keystring>* is the encryption key (in clear text) used for secure communication with all or a specific RADIUS server.

SSH Client Public-Key Authentication

Secure Shell version 2 (SSHv2) is used by ProCurve switches to provide remote access to SSH-enabled management stations. Although SSH provides Telnet-like functions, unlike Telnet, SSH provides encrypted, two-way authenticated transactions. SSH client public-key authentication is one of the types of authentication used.

Client public-key authentication uses one or more public keys (from clients) that must be stored on the switch. Only a client with a private key that matches a public key stored on the switch can gain access at the manager or operator level. For more information about how to configure and use SSH public keys to authenticate SSH clients that try to connect to the switch, refer to “Configuring Secure Shell (SSH)” on page 8-1 in this guide.

The SSH security credential that is stored in the running configuration file is configured with the **ip ssh public-key** command used to authenticate SSH clients for manager or operator access, along with the hashed content of each SSH client public-key.

Syntax: ip ssh public-key <manager loperator> keystring

Set a key for public-key authentication.

manager: *allows manager-level access using SSH public-key authentication.*

operator: *allows operator-level access using SSH public-key authentication.*

"keystring": a legal SSHv2 (RSA or DSA) public key. The text string for the public key must be a single quoted token. If the keystring contains double-quotes, it can be quoted with single quotes ('keystring'). The following restrictions for a keystring apply:

- *A keystring cannot contain both single and double quotes.*
- *A keystring cannot have extra characters, such as a blank space or a new line. However, to improve readability, you can add a backlash at the end of each line.*

Note

The **ip ssh public-key** command allows you to configure only one SSH client public-key at a time. The **ip ssh public-key** command behavior includes an implicit append that never overwrites existing public-key configurations on a running switch.

If you download a software configuration file that contains SSH client public-key configurations, the downloaded public-keys overwrite any existing keys, as happens with any other configured values.

To display the SSH public-key configurations (72 characters per line) stored in a configuration file, enter the **show config** or **show running-config** command. The following example shows the SSH public keys configured for manager access, along with the hashed content of each SSH client public-key, that are stored in a configuration file:

```
...
include-credentials
ip ssh public-key manager "ssh-dss \
AAAAB3NzaC1kc3MAAACBAPwJHSJmTRtpZ9BUNC+ZrsxhMuZEXQhaDME1vc/ \
EvYnTKxQ31bWvr/bT7W58NX/YJ1ZKTV2GZ2QJCicUUZVWjNfJCSa0v03XS4 \
BhkXjtHhz6gD701otgizU006/Xzf4/J9XkJKHkOCnbHIqtB1sbRYBTxj3NzA \
KlymvIaU09X5TDAAAFQCPwKxnbwFfTPasXnxfvDuLSxaC7wAAAIASBwxUP \
pv2scqPPXQghgaTkdPwGGtdFW/+K4xRskAnIaxuG0qLbnekohi+ND4TkKZd \
EeidgDh7qHusBhOFXM2g73RpE2rNqQnSf/QV95kdNwWibxuusBAzvfaJptd \
gca6cYR4xS4TuBcaKiorYj60kk144E1fkDWieQx8zABQAAAIEAu7/1kVOdS \
G0vE0eJD23TLXvu94plXhRKCUAyvv2UyK+piG+Q1ellw9zsMaxPA1XJzSY/ \
imEp4p6WXEMcl0lpXMRnkhnuMMpaPMaQUt8NJTnu6hqf/LdQ2kqZjUuIyV9 \
LWylg5ybS1kFLeOt0oo2Jbpy+U2e4jh2Bb77sX3G5C0= spock@sfc.gov" \
ip ssh public-key manager 'ssh-rsa \
AAAAB3NzaC1yc2EAAAADAQABAAQGDyO9RDD52JZP8k2F2YZXubgwRAN0R \
JRslEov6y1RK3XkmgVatzl+mspiEmPS4wNK7bX/IoXNdGrGkoE8tPklZOZ \
oqGCf5Zs50PlnkXvAidFs55AWqOf4MhfCqvtQCelnt6LFh4ZMig+YewgQG \
M6H1geCSLUBXXSCipdPHysakw== "TectiaClientKey [1024-bit rsa, \
nobody@testmachine, Mon Aug 15 2005 14:47:34]"'
ip ssh public-key manager "ssh-rsa \
AAAAB3NzaC1yc2EAAAABIwAAAIEA1Kk9sVQ9LJOR6XO/hCMPxbiMNOK8C/ay \
+SQ10qGw+K9m3w3TmCfjh0ud9hivgbFT4F99AgnQkvm2eVsgoTtLRnf7uw \
NmpzqOqpHjD9YzItUgSK1uPuFwXMCHKUGKa+G46A+EWxDAIypwVIZ697QmM \
qPFjlzdI4sIo5bDett2d0= joe@hp.com"
...
```

Figure 2-5. Example of SSH Public Keys

If a switch configuration contains multiple SSH client public keys, each public key is saved as a separate entry in the configuration file. You can configure up to ten SSH client public-keys on a switch.

Operating Notes

Caution

- When you first enter the **include-credentials** command to save the additional security credentials to the running configuration, these settings are moved from internal storage on the switch to the running-config file.

You are prompted by a warning message to perform a **write memory** operation to save the security credentials to the startup configuration. The message reminds you that if you do not save the current values of these security settings from the running configuration, they will be lost the next time you boot the switch and will revert to the values stored in the startup configuration.

- When you boot a switch with a startup configuration file that contains the **include-credentials** command, any security credentials that are stored in internal flash memory are ignored and erased. The switch will load only the security settings in the startup configuration file.
- Security settings are no longer automatically saved internally in flash memory and loaded with the startup configuration when a switch boots up. The configuration of all security credentials requires that you use the **write memory** command to save them in the startup configuration in order for them to not be lost when you log off. A warning message reminds you to permanently save a security setting.
- After you enter the **include-credentials** command, the currently configured manager and operator usernames and passwords, RADIUS shared secret keys, SNMP and 802.1X authenticator (port-access) security credentials, and SSH client public-keys are saved in the running configuration.

Use the **no include-credentials** command to disable the display and copying of these security parameters from the running configuration (using the **show running-config** and **copy running-config** commands), without disabling the configured security settings on the switch.

After you enter the **include-credentials** command, you can toggle between the non-display and display of security credentials in **show** and **copy** command output by alternately entering the **no include-credentials** and **include-credentials** commands.

- After you permanently save security configurations to the current startup-config file using the **write memory** command, you can view and manage security settings with the following commands:
 - **show config:** Displays the configuration settings in the current startup-config file.

- **copy config <source-filename> config <target-filename>**: Makes a local copy of an existing startup-config file by copying the contents of the startup-config file in one memory slot to a new startup-config file in another, empty memory slot.
- **copy config tftp**: Uploads a configuration file from the switch to a TFTP server.
- **copy tftp config**: Downloads a configuration file from a TFTP server to the switch.
- **copy config xmodem**: Uploads a configuration file from the switch to an Xmodem host.
- **copy xmodem config**: Downloads a configuration file from an Xmodem host to the switch.

For more information, see “Transferring Startup-Config Files To or From a Remote Server” in the *Management and Configuration Guide*.

- The switch can store up to three configuration files. Each configuration file contains its own security credentials and these security configurations may differ. It is the responsibility of the system administrator to ensure that the appropriate security credentials are contained in the configuration file that is loaded with each software image and that all security credentials in the file are supported.
- If you have already enabled the storage of security credentials (including local manager and operator passwords) by entering the **include-credentials** command, the **Reset-on-clear** option is disabled. When you press the Clear button on the front panel, the manager and operator usernames and passwords are deleted from the running configuration. However, the switch does not reboot after the local passwords are erased. (The **reset-on-clear** option normally reboots the switch when you press the Clear button.)

For more information about the **Reset-on-clear** option and other front-panel security features, see “Configuring Front-Panel Security” on page 2-27 in this guide.

Restrictions

The following restrictions apply when you enable security credentials to be stored in the running configuration with the **include-credentials** command:

- The private keys of an SSH host cannot be stored in the running configuration. Only the public keys used to authenticate SSH clients can be stored. An SSH host's private key is only stored internally, for example, on the switch or on an SSH client device.
- SNMPv3 security credentials saved to a configuration file on a switch cannot be used after downloading the file on a different switch. The SNMPv3 security parameters in the file are only supported when loaded on the same switch for which they were configured. This is because when SNMPv3 security credentials are saved to a configuration file, they are saved with the engine ID of the switch as shown here:

```
snmpv3 engine-id 00:00:00:0b:00:00:08:00:09:01:10:01
```

If you download a configuration file with saved SNMPv3 security credentials on a switch, when the switch loads the file with the current software version the SNMPv3 engine ID value in the downloaded file must match the engine ID of the switch in order for the SNMPv3 users to be configured with the authentication and privacy passwords in the file. (To display the engine ID of a switch, enter the **show snmpv3 engine-id** command. To configure authentication and privacy passwords for SNMPv3 users, enter the **snmpv3 user** command.)

If the engine ID in the saved SNMPv3 security settings in a downloaded configuration file does not match the engine ID of the switch:

- The SNMPv3 users are configured, but without the authentication and privacy passwords. You must manually configure these passwords on the switch before the users can have SNMPv3 access with the privileges you want.
- Only the **snmpv3 user <user_name>** credentials from the SNMPv3 settings in a downloaded configuration file are loaded on the switch, for example:

```
snmpv3 user boris
snmpv3 user alan
```

- You can store 802.1X authenticator (port-access) credentials in a configuration file. However, 802.1X supplicant credentials cannot be stored.
- The local operator password configured with the **password** command is no longer accepted as an 802.1X authenticator credential. A new configuration command (**password port-access**) is introduced to configure

Configuring Username and Password Security

Saving Security Credentials in a Config File

the username and password used as 802.1X authentication credentials for access to the switch. You can store the **password port-access** values in the running configuration file by using the **include-credentials** command.

Note that the **password port-access** values are configured separately from local operator username and passwords configured with the **password operator** command and used for management access to the switch. For more information about how to use the **password port-access** command to configure operator passwords and usernames for 802.1X authentication, see “Do These Steps Before You Configure 802.1X Operation” on page 13-14 in this guide.

Front-Panel Security

The front-panel security features provide the ability to independently enable or disable some of the functions of the two buttons located on the front of the switch for clearing the password (Clear button) or restoring the switch to its factory default configuration (Reset+Clear buttons together). The ability to disable Password Recovery is also provided for situations which require a higher level of switch security.

The front-panel Security features are designed to prevent malicious users from:

- Resetting the password(s) by pressing the Clear button
- Restoring the factory default configuration by using the Reset+Clear button combination.
- Gaining management access to the switch by having physical access to the switch itself

When Security Is Important

Some customers require a high level of security for information. Also, the Health Insurance Portability and Accountability Act (HIPAA) of 1996 requires that systems handling and transmitting confidential medical records must be secure.

It used to be assumed that only system and network administrators would be able to get access to a network switch because switches were typically placed in secure locations under lock and key. For some customers this is no longer true. Others simply want the added assurance that even if someone did manage to get to the switch that data would still remain secure.

If you do not invoke front-panel security on the switch, user-defined passwords can be deleted by pushing the Clear button on the front panel. This function exists so that if customers forget the defined passwords they can still get back into the switch and reset the passwords. This does, however, leave the switch vulnerable when it is located in an area where non-authorized people have access to it. Passwords could easily be cleared by pressing the Clear button. Someone who has physical access to the switch may be able to erase the passwords (and possibly configure new passwords) and take control of the switch.

As a result of increased security concerns, customers now have the ability to stop someone from removing passwords by disabling the Clear and/or Reset buttons on the front of the switch.

Front-Panel Button Functions

This section describes the functionality of the Clear and Reset buttons located on the front panel of the switch.




Figure 2-6. Front-Panel Button Locations on a ProCurve 6120G/XG Switch




Figure 2-7. Front-Panel Button Locations on a ProCurve 6120XG Switch

Clear Button

Pressing the Clear button alone for five seconds resets the password(s) configured on the switch.




Figure 2-8. Press the Clear Button for Five Seconds To Reset the Password(s)

Reset Button

Pressing the Reset button alone for one second causes the switch to reboot.





Figure 2-9. Press and hold the Reset Button for One Second To Reboot the Switch


Restoring the Factory Default Configuration

You can also use the Reset button *together* with the Clear button (Reset+Clear) to **restore the factory default configuration** for the switch. To do this:


1. Press and hold the Reset button.




2. While holding the Reset button, press and hold the Clear button for five seconds.



3. Release the Reset button.



4. If the Clear button is held for greater than 2.5 seconds, configuration will be cleared, and the switch will reboot. It can take approximately 20-25 seconds for the switch to reboot. This process restores the switch configuration to the factory default settings.



Configuring Front-Panel Security

Using the **front-panel-security** command from the global configuration context in the CLI you can:

- Disable or re-enable the password-clearing function of the Clear button. Disabling the Clear button means that pressing it does not remove local password protection from the switch. (This action affects the Clear button when used alone, but does not affect the operation of the Reset+Clear combination described under “Restoring the Factory Default Configuration” on page 2-25.)
- Configure the Clear button to reboot the switch after clearing any local usernames and passwords. This provides an immediate, visual means (plus an Event Log message) for verifying that any usernames and passwords in the switch have been cleared.
- Modify the operation of the Reset+Clear combination (page 2-25) so that the switch still reboots, but does *not* restore the switch’s factory default configuration settings. (Use of the Reset button alone, to simply reboot the switch, is not affected.)
- Disable or re-enable Password Recovery.

Syntax: show front-panel-security

Displays the current front-panel-security settings:

Clear Password: *Shows the status of the Clear button on the front panel of the switch. **Enabled** means that pressing the Clear button erases the local usernames and passwords configured on the switch (and thus removes local password protection from the switch). **Disabled** means that pressing the Clear button does not remove the local usernames and passwords configured on the switch. (Default: **Enabled**.)*

Reset-on-clear: Shows the status of the reset-on-clear option (**Enabled** or **Disabled**). When reset-on-clear is disabled and Clear Password is enabled, then pressing the Clear button erases the local usernames and passwords from the switch. When reset-on-clear is enabled, pressing the Clear button erases the local usernames and passwords from the switch and reboots the switch. (Enabling **reset-on-clear** automatically enables **clear-password**.) (Default: **Disabled**.)

Note: If you have stored security credentials (including the local manager and operator usernames and passwords) to the running config file by entering the **include-credentials** command, the **Reset-on-clear** option is ignored. If you press the Clear button on the front panel, the manager and operator usernames and passwords are deleted from the startup configuration file, but the switch does not reboot. For more information about storing security credentials, see “Saving Security Credentials in a Config File” on page 2-10 in this guide.

Factory Reset: Shows the status of the System Reset button on the front panel of the switch. Enabled means that pressing the System Reset button reboots the switch and also enables the System Reset button to be used with the Clear button (page 2-25) to reset the switch to its factory-default configuration. (Default: **Enabled**.)

Password Recovery: Shows whether the switch is configured with the ability to recover a lost password. (Refer to “Password Recovery Process” on page 2-34.) (Default: **Enabled**.)

CAUTION: Disabling this option removes the ability to recover a password on the switch. Disabling this option is an extreme measure and is not recommended unless you have the most urgent need for high security. If you disable password-recovery and then lose the password, you will have to use the Reset and Clear buttons (page 2-25) to reset the switch to its factory-default configuration and create a new password.

For example, **show front-panel-security** produces the following output when the switch is configured with the default front-panel security settings.

```
ProCurve(config)# show front-panel-security
Clear Password      - Enabled
Reset-on-clear     - Disabled
Factory Reset      - Enabled
Password Recovery   - Enabled
```

Figure 2-10. The Default Front-Panel Security Settings

Disabling the Clear Password Function of the Clear Button

Syntax: no front-panel-security password-clear

In the factory-default configuration, pressing the Clear button on the switch's front panel erases any local usernames and passwords configured on the switch. This command disables the password clear function of the Clear button, so that pressing it has no effect on any local usernames and passwords.

*(Default: **Enabled**.)*

Note: Although the Clear button does not erase passwords when disabled, you can still use it with the Reset button (Reset+Clear) to restore the switch to its factory default configuration, as described under “Restoring the Factory Default Configuration” on page 2-25.

This command displays a Caution message in the CLI. If you want to proceed with disabling the Clear button, type **[Y]**; otherwise type **[N]**. For example:

```
ProCurve(config)# no front-panel-security password-clear
**** CAUTION ****
Disabling the clear button prevents switch passwords from being easily reset or
recovered. Ensure that you are familiar with the front-panel security options
before proceeding.
Continue with disabling the clear button [y/n]? y
ProCurve(config)# show front-panel-security
Clear Password      - Disabled
Factory Reset       - Enabled
Password Recovery   - Enabled
```

Indicates the command has disabled the Clear button on the switch's front panel. In this case the Show command does not include the **reset-on-clear** status because it is inoperable while the Clear Password functionality is disabled, and must be reconfigured whenever Clear Password is re-enabled.

Re-Enabling the Clear Button and Setting or Changing the “Reset-On-Clear” Operation

Syntax: [no] front-panel-security password-clear reset-on-clear

This command does both of the following:

- *Re-enables the password-clearing function of the Clear button on the switch’s front panel.*
- *Specifies whether the switch reboots if the Clear button is pressed.*

*To re-enable password-clear, you must also specify whether to enable or disable the **reset-on-clear** option.*

Defaults:

- password-clear: **Enabled**.
- reset-on-clear: **Disabled**.

Thus:

- *To enable password-clear with reset-on-clear disabled, use this syntax:*

no front-panel-security password-clear reset-on-clear

- *To enable password-clear with reset-on-clear also enabled, use this syntax:*

front-panel-security password-clear reset-on-clear

(Either form of the command enables password-clear.)

Note: *If you disable **password-clear** and also disable the **password-recovery** option, you can still recover from a lost password by using the Reset+Clear button combination at reboot as described on page 2-25. Although the Clear button does not erase passwords when disabled, you can still use it with the Reset button (Reset+Clear) to restore the switch to its factory default configuration. You can then get access to the switch to set a new password.*

For example, suppose that **password-clear** is disabled and you want to restore it to its default configuration (enabled, with **reset-on-clear** disabled).




Figure 2-12. Example of Re-Enabling the Clear Button’s Default Operation

Changing the Operation of the Reset+Clear Combination

In their default configuration, using the Reset+Clear buttons in the combination described under “Restoring the Factory Default Configuration” on page 2-25 replaces the switch’s current startup-config file with the factory-default startup-config file, then reboots the switch, and removes local password protection. *This means that anyone who has physical access to the switch could use this button combination to replace the switch’s current configuration with the factory-default configuration, and render the switch accessible without the need to input a username or password.* You can use the **factory-reset** command to prevent the Reset+Clear combination from being used for this purpose.

Syntax: [no] front-panel-security factory-reset

Disables or re-enables the following functions associated with using the Reset+Clear buttons in the combination described under “Restoring the Factory Default Configuration” on page 2-25:

- *Replacing the current startup-config file with the factory-default startup-config file*
- *Clearing any local usernames and passwords configured on the switch*

(Default: Both functions enabled.)

Notes: The Reset+Clear button combination always reboots the switch, regardless of whether the “no” form of the command has been used to disable the above two functions. Also, if you disable **factory-reset**, you cannot disable the **password-recovery** option, and the reverse.

```
ProCurve(config)# no front-panel-security factory-reset
```

**** CAUTION ****

Disabling the factory reset option prevents switch configuration and passwords from being easily reset or recovered. Ensure that you are familiar with the front panel security options before proceeding.

Continue with disabling the factory reset option[y/n]? y

```
ProCurve(config)# show front-panel-security
```

Clear Password	- Enabled
Reset-on-clear	- Disabled
Factory Reset	- Disabled
Password Recovery	- Enabled

The command to disable the factory-reset operation produces this caution. To complete the command, press [Y]. To abort the command, press [N].

Completes the command to disable the factory reset option.

Displays the current front-panel-security configuration, with Factory Reset disabled.

Figure 2-13. Example of Disabling the Factory Reset Option

Password Recovery

The password recovery feature is enabled by default and provides a method for regaining management access to the switch (without resetting the switch to its factory default configuration) in the event that the system administrator loses the local manager username (if configured) or password. Using Password Recovery requires:

- **password-recovery** enabled (the default) on the switch prior to an attempt to recover from a lost username/password situation
- Contacting your ProCurve Customer Care Center to acquire a one-time-use password

Disabling or Re-Enabling the Password Recovery Process

Disabling the password recovery process means that the only method for recovering from a lost manager username (if configured) and password is to reset the switch to its factory-default configuration, which removes any non-default configuration settings.

Caution

Disabling **password-recovery** requires that **factory-reset** be enabled, and locks out the ability to recover a lost manager username (if configured) and password on the switch. In this event, there is no way to recover from a lost manager username/password situation without resetting the switch to its

factory-default configuration. This can disrupt network operation and make it necessary to temporarily disconnect the switch from the network to prevent unauthorized access and other problems while it is being reconfigured. Also, with **factory-reset** enabled, unauthorized users can use the Reset+Clear button combination to reset the switch to factory-default configuration and gain management access to the switch.

Syntax: [no] front-panel-security password-recovery

Enables or (using the “no” form of the command) disables the ability to recover a lost password.

When this feature is enabled, the switch allows management access through the password recovery process described below. This provides a method for recovering from a lost manager username (if configured) and password. When this feature is disabled, the password recovery process is disabled and the only way to regain management access to the switch is to use the Reset+Clear button combination (page 2-25) to restore the switch to its factory default configuration.

Note: To disable **password-recovery**:

- You must have physical access to the front panel of the switch.
- The **factory-reset** parameter must be enabled (the default).

(Default: Enabled.)

Steps for Disabling Password-Recovery.

1. Set the CLI to the global interface context.
2. Use **show front-panel-security** to determine whether the factory-reset parameter is enabled. If it is disabled, use the **front-panel-security factory-reset** command to enable it.
3. Press and release the Clear button on the front panel of the switch.
4. Within 60-seconds of pressing the Clear button, enter the following command:

no front-panel-security password-recovery

5. Do one of the following after the “**CAUTION**” message appears:
 - If you want to complete the command, press **[Y]** (for “Yes”).
 - If you want to abort the command, press **[N]** (for “No”).

Figure 2-14 shows an example of disabling the **password-recovery** parameter.

```
ProCurve(config)# no front-panel-security password-recovery
                        **** CAUTION ****
Disabling the clear button without password recovery prevents switch passwords
from being reset.  If the switch password is lost, restoring the default factory
configuration will be required to regain access!

Continue with disabling password recovery [y/n]? y

ProCurve(config)# _
```

Figure 2-14. Example of the Steps for Disabling Password-Recovery

Password Recovery Process

If you have lost the switch's manager username/password, but **password-recovery** is enabled, then you can use the Password Recovery Process to gain management access to the switch with an alternate password supplied by ProCurve.

Note

If you have disabled **password-recovery**, which locks out the ability to recover a manager username/password pair on the switch, then the only way to recover from a lost manager username/password pair is to use the Reset+Clear button combination described under “Restoring the Factory Default Configuration” on page 2-25. This can disrupt network operation and make it necessary to temporarily disconnect the switch from the network to prevent unauthorized access and other problems while it is being reconfigured.

To use the **password-recovery** option to recover a lost password:

1. Note the switch's base MAC address. It is shown on the label located on the upper right front corner of the switch.
2. Contact your ProCurve Customer Care Center for further assistance. Using the switch's MAC address, the ProCurve Customer Care Center will generate and provide a “one-time use” alternate password you can use with the to gain management access to the switch. Once you gain access, you can configure a new, known password.

Note

The alternate password provided by the ProCurve Customer Care Center is valid only for a single login attempt. You cannot use the *same “one-time-use” password* if you lose the password a second time. Because the password algorithm is randomized based upon your switch's MAC address, the password will change as soon as you use the “one-time-use” password provided to you by the ProCurve Customer Care Center.

Web and MAC Authentication

Contents

Overview	3-3
Web Authentication	3-3
MAC Authentication	3-4
Concurrent Web and MAC Authentication	3-4
Authorized and Unauthorized Client VLANs	3-5
RADIUS-Based Authentication	3-6
Wireless Clients	3-6
How Web and MAC Authentication Operate	3-6
Web-based Authentication	3-7
MAC-based Authentication	3-9
Terminology	3-11
Operating Rules and Notes	3-12
Setup Procedure for Web/MAC Authentication	3-14
Before You Configure Web/MAC Authentication	3-14
Configuring the RADIUS Server To Support MAC Authentication	3-16
Configuring the Switch To Access a RADIUS Server	3-17
Configuring Web Authentication	3-20
Overview	3-20
Configuration Commands for Web Authentication	3-21
Show Commands for Web Authentication	3-28
Customizing Web Authentication HTML Files (Optional)	3-34
Implementing Customized Web-Auth Pages	3-34
Operating Notes and Guidelines	3-34
Customizing HTML Templates	3-35
Customizable HTML Templates	3-36
Configuring MAC Authentication on the Switch	3-50

Overview	3-50
Configuration Commands for MAC Authentication	3-51
Show Commands for MAC-Based Authentication	3-54
Client Status	3-60

Overview

Feature	Default	Menu	CLI	Web
Configure Web Authentication	n/a	—	3-20	—
Configure MAC Authentication	n/a	—	3-50	—
Display Web Authentication Status and Configuration	n/a	—	3-28	—
Display MAC Authentication Status and Configuration	n/a	—	3-54	—

Web and MAC authentication are designed for employment on the “edge” of a network to provide port-based security measures for protecting private networks and a switch from unauthorized access. Because neither method requires clients to run special supplicant software (unlike 802.1X authentication), both Web and MAC authentication are suitable for legacy systems and temporary access situations where introducing supplicant software is not an attractive option. Only a web browser (for Web authentication) or a MAC address (for MAC authentication) is required.

Both Web and MAC authentication methods rely on a RADIUS server to authenticate network access. This simplifies access security management by allowing you to control access from a master database in a single server. (You can use up to three RADIUS servers to provide backups in case access to the primary server fails.) It also means the same credentials can be used for authentication, regardless of which switch or switch port is the current access point into the LAN.

On a port configured for Web or MAC Authentication, the switch operates as a port-access authenticator using a RADIUS server and the CHAP protocol. Inbound traffic is processed by the switch alone, until authentication occurs. Some traffic from the switch to an unauthorized client is supported (for example, broadcast or unknown destination packets) before authentication occurs.

Web Authentication

The Web Authentication (Web-Auth) method uses a web page login to authenticate users for access to the network. When a client connects to the switch and opens a web browser, the switch automatically presents a login page.

Note

A proxy server is not supported for use by a browser on a client device that accesses the network through a port configured for web authentication.

- In the login page, a client enters a username and password, which the switch forwards to a RADIUS server for authentication. After authenticating a client, the switch grants access to the secured network. Besides a web browser, the client needs no special supplicant software.

MAC Authentication

The MAC Authentication (MAC-Auth) method grants access to a secure network by authenticating devices for access to the network. When a device connects to the switch, either by direct link or through the network, the switch forwards the device's MAC address to the RADIUS server for authentication. The RADIUS server uses the device MAC address as the username and password, and grants or denies network access in the same way that it does for clients capable of interactive logons. (The process does not use either a client device configuration or a logon session.) MAC authentication is well-suited for clients that are not capable of providing interactive logons, such as telephones, printers, and wireless access points. Also, because most RADIUS servers allow for authentication to depend on the source switch and port through which the client connects to the network, you can use MAC-Auth to “lock” a particular device to a specific switch and port.

Note

802.1X port-access, Web authentication, and MAC authentication can be configured at the same time on the same port. A maximum of 32 clients is supported on the port. (The default is one client.)

Web and/or MAC authentication and MAC lockdown, MAC lockout, and port-security are mutually exclusive on a given port. If you configure any of these authentication methods on a port, you must disable LACP on the port.

Concurrent Web and MAC Authentication

Web authentication and MAC authentication can be configured at the same time on a port. It is assumed that MAC authentication will use an existing MAC address. The following conditions apply for concurrent Web and MAC authentication:

- A specific MAC address cannot be authenticated by both Web and MAC authentication at the same time.

- Each new Web/MAC Auth client always initiates a MAC authentication attempt. This same client can also initiate Web authentication at any time before the MAC authentication succeeds. If either authentication succeeds then the other authentication (if in progress) is ended. No further Web/MAC authentication attempts are allowed until the client is deauthenticated.
- Web and MAC authentications are not allowed on the same port if unauthenticated VLAN (that is, a guest VLAN) is enabled for MAC authentication. An unauthenticated VLAN can't be enabled for MAC authentication if Web and MAC authentication are both enabled on the port.
- Hitless reauthentication must be of the same type (MAC) that was used for the initial authentication. Non-hitless reauthentication can be of any type.

The remaining Web/MAC functionality, including interactions with 802.1X, remains the same. Web and MAC authentication can be used for different clients on the same port.

Normally, MAC authentication finishes much sooner than Web authentication. However, if Web authentication should complete first, MAC authentication will cease even though it is possible that MAC authentication could succeed. There is no guarantee that MAC authentication ends before Web authentication begins for the client.

Concurrent Web and MAC authentication is backward compatible with all existing user configurations.

Authorized and Unauthorized Client VLANs

Web-Auth and MAC-Auth provide a port-based solution in which a port belongs to one, untagged VLAN at a time. The switch supports up to 32 simultaneous client sessions per port. All authenticated client sessions operate in the same untagged VLAN. (If you want the switch to simultaneously support multiple client sessions in different VLANs for a network application, design your system so that clients request network access on different switch ports.)

In the default configuration, the switch blocks access to all clients that the RADIUS server does not authenticate. However, you can configure an individual port to provide limited network services and access to unauthorized

clients by using an “unauthorized” VLAN for each session. The unauthorized VLAN ID assignment can be the same for all ports, or different, depending on the services and access you plan to allow for unauthenticated clients.

You configure access to an optional, unauthorized VLAN when you configure Web and MAC authentication on a port.

RADIUS-Based Authentication

In Web and MAC authentication, you use a RADIUS server to temporarily assign a port to a static VLAN to support an authenticated client. When a RADIUS server authenticates a client, the switch-port membership during the client’s connection is determined according to the following hierarchy:

1. A RADIUS-assigned VLAN
2. An authorized VLAN specified in the Web- or MAC-Auth configuration for the subject port.
3. A static, port-based, untagged VLAN to which the port is configured. A RADIUS-assigned VLAN has priority over switch-port membership in any VLAN.

Wireless Clients

You can allow wireless clients to move between switch ports under Web/MAC Authentication control. Clients may move from one Web-authorized port to another or from one MAC-authorized port to another. This capability allows wireless clients to move from one access point to another without having to reauthenticate.


How Web and MAC Authentication Operate

Before gaining access to the network, a client first presents authentication credentials to the switch. The switch then verifies the credentials with a RADIUS authentication server. Successfully authenticated clients receive access to the network, as defined by the System Administrator. Clients who fail to authenticate successfully receive no network access or limited network access as defined by the System Administrator.

Web-based Authentication

When a client connects to a Web-Auth enabled port, communication is redirected to the switch. A temporary IP address is assigned by the switch and a login screen is presented for the client to enter their username and password.

The default User Login screen is shown in Figure 3-1.



User Login

In order to access this network, you must first log in.

Username:

Password:

Figure 3-1. Example of Default User Login Screen

When a client connects to the switch, it sends a DHCP request to receive an IP address to connect to the network. To avoid address conflicts in a secure network, you can specify a temporary IP address pool to be used by DHCP by configuring the **dhcp-addr** and **dhcp-lease** options when you enable web authentication with the **aaa port-access web-based** command.

The Secure Socket Layer (SSLv3/TLSv1) feature provides remote web access to the network via authenticated transactions and encrypted paths between the switch and management station clients capable of SSL/TLS. If you have enabled SSL on the switch, you can specify the **ssl-login** option when you configure web authentication so that clients who log in to specified ports are redirected to a secure login page (<https://...>) to enter their credentials.

The switch passes the supplied username and password to the RADIUS server for authentication and displays the following progress message:



Authenticating...

Please wait while your credentials are verified.

Figure 3-2. Progress Message During Authentication

If the client is authenticated and the maximum number of clients allowed on the port (**client-limit**) has not been reached, the port is assigned to a static, untagged VLAN for network access. After a successful login, a client may be redirected to a URL if you specify a URL value (**redirect-url**) when you configure web authentication.



Figure 3-3. Authentication Completed

The assigned VLAN is determined, in order of priority, as follows:

1. If there is a RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to this VLAN and temporarily drops all other VLAN memberships.
2. If there is no RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to the authorized VLAN (**auth-vid** if configured) and temporarily drops all other VLAN memberships.
3. If neither 1 or 2, above, apply, but the port is an untagged member of a statically configured, port-based VLAN, then the port remains in this VLAN.
4. If neither 1, 2, or 3, above, apply, then the client session does not have access to any statically configured, untagged VLANs and client access is blocked.

The assigned port VLAN remains in place until the session ends. Clients may be forced to reauthenticate after a fixed period of time (**reauth-period**) or at any time during a session (**reauthenticate**). An implicit logoff period can be set if there is no activity from the client after a given amount of time (**logoff-period**). In addition, a session ends if the link on the port is lost, requiring reauthentication of all clients. Also, if a client moves from one port to another and client moves have not been enabled (**client-moves**) on the ports, the session ends and the client must reauthenticate for network access. At the end of the session the port returns to its pre-authentication state. Any changes to the port's VLAN memberships made while it is an authorized port take affect at the end of the session.

A client may not be authenticated due to invalid credentials or a RADIUS server timeout. The **max-retries** parameter specifies how many times a client may enter their credentials before authentication fails. The **server-timeout** parameter sets how long the switch waits to receive a response from the RADIUS server before timing out. The **max-requests** parameter specifies how many authentication attempts may result in a RADIUS server timeout before authentication fails. The switch waits a specified amount of time (**quiet-period**) before processing any new authentication requests from the client.

Network administrators may assign unauthenticated clients to a specific static, untagged VLAN (**unauth-vid**), to provide access to specific (guest) network resources. If no VLAN is assigned to unauthenticated clients the port is blocked and no network access is available. Should another client successfully authenticate through that port any unauthenticated clients on the **unauth-vid** are dropped from the port.

MAC-based Authentication

When a client connects to a MAC-Auth enabled port traffic is blocked. The switch immediately submits the client's MAC address (in the format specified by the **addr-format**) as its certification credentials to the RADIUS server for authentication.

If the client is authenticated and the maximum number of MAC addresses allowed on the port (**addr-limit**) has not been reached, the port is assigned to a static, untagged VLAN for network access.

The assigned VLAN is determined, in order of priority, as follows:

1. If there is a RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to this VLAN and temporarily drops all other VLAN memberships.
2. If there is no RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to the Authorized VLAN (**auth-vid** if configured) and temporarily drops all other VLAN memberships.
3. If neither 1 or 2, above, apply, but the port is an untagged member of a statically configured, port-based VLAN, then the port remains in this VLAN.
4. If neither 1, 2, or 3, above, apply, then the client session does not have access to any statically configured, untagged VLANs and client access is blocked.

The assigned port VLAN remains in place until the session ends. Clients may be forced to reauthenticate after a fixed period of time (**reauth-period**) or at any time during a session (**reauthenticate**). An implicit logoff period can be set if there is no activity from the client after a given amount of time (**logoff-period**). In addition, a session ends if the link on the port is lost, requiring reauthentication of all clients. Also, if a client moves from one port to another and client moves have not been enabled (**addr-moves**) on the ports, the session ends and the client must reauthenticate for network access. At the end of the session the port returns to its pre-authentication state. Any changes to the port's VLAN memberships made while it is an authenticated port take affect at the end of the session.

A client may not be authenticated due to invalid credentials or a RADIUS server timeout. The **server-timeout** parameter sets how long the switch waits to receive a response from the RADIUS server before timing out. The **max-requests** parameter specifies how many authentication attempts may result in a RADIUS server timeout before authentication fails. The switch waits a specified amount of time (**quiet-period**) before processing any new authentication requests from the client.

Network administrators may assign unauthenticated clients to a specific static, untagged VLAN (**unauth-vid**), to provide access to specific (guest) network resources. If no VLAN is assigned to unauthenticated clients the port remains in its original VLAN configuration. Should another client successfully authenticate through that port any unauthenticated clients are dropped from the port.

Terminology

Authorized-Client VLAN: Like the Unauthorized-Client VLAN, this is a conventional, static, untagged, port-based VLAN previously configured on the switch by the System Administrator. The intent in using this VLAN is to provide authenticated clients with network access and services. When the client connection terminates, the port drops its membership in this VLAN.

Authentication Server: The entity providing an authentication service to the switch.

Authenticator: In ProCurve switch applications, a device such as a ProCurve Switch 8212zl that requires a client or device to provide the proper credentials (MAC address, or username and password) before being allowed access to the network.

CHAP: Challenge Handshake Authentication Protocol. Also known as “CHAP-RADIUS”.

Client: In this application, an end-node device such as a management station, workstation, or mobile PC linked to the switch through a point-to-point LAN link.

Redirect URL: A System Administrator-specified web page presented to an authorized client following Web Authentication. ProCurve recommends specifying this URL when configuring Web Authentication on a switch. Refer to **aaa port-access web-based [e] < port-list > [redirect-url < url >]** on page 3-27.

Static VLAN: A VLAN that has been configured as “permanent” on the switch by using the CLI **vlan < vid >** command or the Menu interface.

Unauthorized-Client VLAN: A conventional, static, untagged, port-based VLAN previously configured on the switch by the System Administrator. It is used to provide limited network access and services to clients who are not authenticated.

Operating Rules and Notes

- The switch supports concurrent 802.1X, Web and MAC authentication operation on a port (with up to 2 clients allowed). However, concurrent operation of Web and MAC authentication with other types of authentication on the same port is not supported. That is, the following authentication types are *mutually exclusive* on a given port:
 - Web and/or MAC Authentication (with or without 802.1X)
 - MAC lockdown
 - MAC lockout
 - Port-Security
- Order of Precedence for Port Access Management (highest to lowest):
 - a. MAC lockout
 - b. MAC lockdown or Port Security
 - c. Port-based Access Control (802.1X) or Web Authentication or MAC Authentication

Port Access Management

When configuring a port for Web or MAC Authentication, be sure that a higher precedent port access management feature is not enabled on the port. For example, be sure that Port Security is disabled on a port before configuring the port for Web or MAC Authentication. If Port Security is enabled on the port this misconfiguration does not allow Web or MAC Authentication to occur.

- VLANs: If your LAN does not use multiple VLANs, then you do not need to configure VLAN assignments in your RADIUS server or consider using either Authorized or Unauthorized VLANs. If your LAN does use multiple VLANs, then some of the following factors may apply to your use of Web-Auth and MAC-Auth.
 - Web-Auth and MAC-Auth operate only with port-based VLANs. Operation with protocol VLANs is not supported, and clients do not have access to protocol VLANs during Web-Auth and MAC-Auth sessions.
 - A port can belong to one, untagged VLAN during any client session. Where multiple authenticated clients may simultaneously use the same port, they must all be capable of operating on the same VLAN.
 - During an authenticated client session, the following hierarchy determines a port's VLAN membership:

1. If there is a RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to this VLAN and temporarily drops all other VLAN memberships.
 2. If there is no RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to the Authorized VLAN (if configured) and temporarily drops all other VLAN memberships.
 3. If neither 1 or 2, above, apply, but the port is an untagged member of a statically configured, port-based VLAN, then the port remains in this VLAN.
 4. If neither 1, 2, or 3, above, apply, then the client session does not have access to any statically configured, untagged VLANs and client access is blocked.
- After an authorized client session begins on a given port, the port's VLAN membership does not change. If other clients on the same port become authenticated with a different VLAN assignment than the first client, the port blocks access to these other clients until the first client session ends.
 - The optional “authorized” VLAN (**auth-vid**) and “unauthorized” VLAN (**unauth-vid**) you can configure for Web- or MAC-based authentication must be statically configured VLANs on the switch. Also, if you configure one or both of these options, any services you want clients in either category to access must be available on those VLANs.
- Where a given port's configuration includes an unauthorized client VLAN assignment, the port will allow an unauthenticated client session only while there are no requests for an authenticated client session on that port. In this case, if there is a successful request for authentication from an authorized client, the switch terminates the unauthorized-client session and begins the authorized-client session.
 - When a port on the switch is configured for Web or MAC Authentication and is supporting a current session with another device, rebooting the switch invokes a re-authentication of the connection.
 - When a port on the switch is configured as a Web- or MAC-based authenticator, it blocks access to a client that does not provide the proper authentication credentials. If the port configuration includes an optional, unauthorized VLAN (**unauth-vid**), the port is temporarily placed in the unauthorized VLAN if there are no other authorized clients currently using the port with a different VLAN assignment. If an authorized client is using the port with a different VLAN or if there is no unauthorized VLAN configured, the unauthorized client does not receive access to the network.
 - Web- or MAC-based authentication and LACP cannot both be enabled on the same port.

Web/MAC Authentication and LACP

Web or MAC authentication and LACP are not supported at the same time on a port. The switch automatically disables LACP on ports configured for Web or MAC authentication.

- Use the **show port-access web-based** commands to display session status, port-access configuration settings, and statistics for Web-Auth sessions.
- When spanning tree is enabled on a switch that uses 802.1X, Web authentication, or MAC authentication, loops may go undetected. For example, spanning tree packets that are looped back to an edge port will not be processed because they have a different broadcast/multicast MAC address from the client-authenticated MAC address. To ensure that client-authenticated edge ports get blocked when loops occur, you should enable loop protection on those ports. For more information, refer to “Loop Protection” in the chapter titled “Multiple Instance Spanning-Tree Operation” in the *Advanced Traffic Management Guide* for your switch.

Setup Procedure for Web/MAC Authentication

Before You Configure Web/MAC Authentication

1. Configure a local username and password on the switch for both the Operator (login) and Manager (enable) access levels. (While this is not required for a Web- or MAC-based configuration, ProCurve recommends that you use a local user name and password pair, at least until your other security measures are in place, to protect the switch configuration from unauthorized access.)
2. Determine the switch ports that you want to configure as authenticators. Note that before you configure Web- or MAC-based authentication on a port operating in an LACP trunk, you must remove the port from the trunk. (For more information, refer to the “Web/MAC Authentication and LACP” on page 3-14.)

To display the current configuration of 802.1X, Web-based, and MAC authentication on all switch ports, enter the **show port-access config** command.

```
ProCurve (config)# show port-access config

Port Access Status Summary
Port-access authenticator activated [No] : Yes
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : Yes

      Supplicant Authenticator Web Auth Mac Auth
Port Enabled      Enabled      Enabled  Enabled
-----
1      Yes         No           No       Yes
2      No          Yes          No       Yes
3      No          Yes          No       No
4      No          No           No       No
5      No          No           No       No
6      No          No           No       No
7      No          No           No       No
8      No          No           No       No
9      No          No           No       No
10     No          No           No       No
11     No          No           No       No
12     No          No           No       No
...

```

Figure 3-4. Example of show port-access config Command Output

3. Determine whether any VLAN assignments are needed for authenticated clients.
 - a. If you configure the RADIUS server to assign a VLAN for an authenticated client, this assignment overrides any VLAN assignments configured on the switch while the authenticated client session remains active. Note that the VLAN must be statically configured on the switch.
 - b. If there is no RADIUS-assigned VLAN, the port can join an “Authorized VLAN” for the duration of the client session, if you choose to configure one. This must be a port-based, statically configured VLAN on the switch.
 - c. If there is neither a RADIUS-assigned VLAN or an “Authorized VLAN” for an authenticated client session on a port, then the port’s VLAN membership remains unchanged during authenticated client sessions. In this case, configure the port for the VLAN in which you want it to operate during client sessions.

Note that when configuring a RADIUS server to assign a VLAN, you can use either the VLAN's name or VID. For example, if a VLAN configured in the switch has a VID of 100 and is named **vlan100**, you could configure the RADIUS server to use either "100" or "vlan100" to specify the VLAN.

4. Determine whether to use the optional "Unauthorized VLAN" mode for clients that the RADIUS server does not authenticate. This VLAN must be statically configured on the switch. If you do not configure an "Unauthorized VLAN", the switch simply blocks access to unauthenticated clients trying to use the port.
5. Determine the authentication policy you want on the RADIUS server and configure the server. Refer to the documentation provided with your RADIUS application and include the following in the policy for each client or client device:
 - The CHAP-RADIUS authentication method.
 - An encryption key
 - One of the following:
 - If you are configuring Web-based authentication, include the user name and password for each authorized client.
 - If you are configuring MAC-based authentication, enter the device MAC address in both the username and password fields of the RADIUS policy configuration for that device. Also, if you want to allow a particular device to receive authentication only through a designated port and switch, include this in your policy.
6. Determine the IP address of the RADIUS server(s) you will use to support Web- or MAC-based authentication. (For information on configuring the switch to access RADIUS servers, refer to "Configuring the Switch To Access a RADIUS Server" on page 3-17.)

Configuring the RADIUS Server To Support MAC Authentication

On the RADIUS server, configure the client device authentication in the same way that you would any other client, except:

- Configure the client device's (hexadecimal) MAC address as both username and password. Be careful to configure the switch to use the same format that the RADIUS server uses. Otherwise, the server will deny access. The switch provides four format options:

aabbccddeeff (the default format)

aabbcc-ddeeff

aa-bb-cc-dd-ee-ff
aa:bb:cc:dd:ee:ff
AABBCCDDEEFF
AABBCC-DDEEFF
AA-BB-CC-DD-EE-FF
AA:BB:CC:DD:EE:FF

- If the device is a switch or other VLAN-capable device, use the base MAC address assigned to the device, and not the MAC address assigned to the VLAN through which the device communicates with the authenticator switch. Note that the switch applies a single MAC address to all VLANs configured in the switch. Thus, for a given switch, the MAC address is the same for all VLANs configured on the switch. (Refer to the chapter titled “Static Virtual LANs (VLANs)” in the *Advanced Traffic Management Guide* for your switch.)

Configuring the Switch To Access a RADIUS Server

RADIUS Server Configuration Commands		
radius-server		
[host <ip-address>]		below
[key < global-key-string >]		below
radius-server host <ip-address> [oobm] key <server-specific key-string> [no] radius-server host <ip-address> key		3-18

This section describes the minimal commands for configuring a RADIUS server to support Web-Auth and MAC Auth. For information on other RADIUS command options, refer to chapter 5, “RADIUS Authentication, Authorization, and Accounting” .

Syntax: [no] radius-server

[host < ip-address >] [oobm]

*Adds a server to the RADIUS configuration or (with **no**) deletes a server from the configuration. You can configure up to three RADIUS server addresses. The switch uses the first server it successfully accesses. (Refer to “RADIUS Authentication, Authorization, and Accounting” on page 5-1.)*

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the RADIUS traffic will go through the out-of-band management (OOBM) port.*

[key < global-key-string >]

Specifies the global encryption key the switch uses with servers for which the switch does not have a server-specific key assignment (below). This key is optional if all RADIUS server addresses configured in the switch include a server-specific encryption key. (Default: Null.)

Syntax: radius-server host < ip-address > key <server-specific key-string>
[no] radius-server host < ip-address > key

Optional. Specifies an encryption key for use during authentication (or accounting) sessions with the specified server. This key must match the encryption key used on the RADIUS server. Use this command only if the specified server requires a different encryption key than configured for the global encryption key, above.

*The **no** form of the command removes the key configured for a specific server.*

For example, to configure the switch to access a RADIUS server at IP address 192.168.32.11 using a server specific shared secret key of '1A7rd'

```
ProCurve(config)# radius-server host 192.168.32 11
ProCurve(config)# radius-server host 192.168.32.11 key 1A7rd

ProCurve(config)# show radius

Status and Counters - General RADIUS Information

Deadtime(min) : 0
Timeout(secs) : 5
Retransmit Attempts : 3
Global Encryption Key :

Server IP Addr  Auth  Acct
Port  Port  Encryption Key
-----
192.168.32.11  1812  1813  1A7rd
```

Figure 3-5. Example of Configuring a Switch To Access a RADIUS Server

Configuring Web Authentication

Overview

1. If you have not already done so, configure a local username and password pair on the switch.
2. Identify or create a redirect URL for use by authenticated clients. ProCurve recommends that you provide a redirect URL when using Web Authentication. If a redirect URL is not specified, web browser behavior following authentication may not be acceptable.
3. If you plan to use multiple VLANs with Web Authentication, ensure that these VLANs are configured on the switch and that the appropriate port assignments have been made. Also, confirm that the VLAN used by authorized clients can access the redirect URL.
4. Use the **ping** command in the switch console interface to ensure that the switch can communicate with the RADIUS server you have configured to support Web-Auth on the switch.
5. Configure the switch with the correct IP address and encryption key to access the RADIUS server.
6. (Optional) To use SSL encryption for web authentication login, configure and enable SSL on the switch.
7. Enable web authentication on the switch ports you want to use.
8. Configure the optional settings that you want to use for web authentication; for example:
 - To avoid address conflicts in a secure network, configure the base IP address and mask to be used by the switch for temporary DHCP addresses. You can also set the lease length for these temporary IP addresses.
 - To use SSL encryption for web authentication login, configure the SSL option.
 - To redirect authorized clients to a specified URL, configure the Redirect URL option.
9. Configure how web-authenticator ports transmit traffic before they successfully authenticate a client and enter the authenticated state:
 - You can block incoming and outgoing traffic on a port before authentication occurs.

- You can block only incoming traffic on a port before authentication occurs. Outgoing traffic with unknown destination addresses is flooded on unauthenticated ports configured for web authentication. For example, Wake-on-LAN traffic is transmitted on a web-authenticated egress port that has not yet transitioned to the authenticated state;
10. Test both authorized and unauthorized access to your system to ensure that Web Authentication works properly on the ports you have configured for port-access using Web Authentication.

Note Client web browsers may not use a proxy server to access the network.

Configuration Commands for Web Authentication

Command	Page
Configuration Level	
aaa port-access <port-list> controlled-directions <both in>	3-22
[no] aaa port-access web-based <port-list>	3-24
[auth-vid]	3-24
[clear-statistics]	3-24
[client-limit]	3-24
[dhcp-addr]	3-24
[dhcp-lease]	3-25
[ewa-server]	3-25
[logoff-period]	3-26
[max-requests]	3-26
[max-retries]	3-26
[quiet-period]	3-26
[reauth-period]	3-27
[reauthenticate]	3-27
[redirect-url]	3-27
[server-timeout]	3-27
[unauth-vid]	3-54

Syntax: `aaa port-access <port-list> controlled-directions <both | in>`

*After you enable web-based authentication on specified ports, you can use the **aaa port-access controlled-directions** command to configure how a port transmits traffic before it successfully authenticates a client and enters the authenticated state.*

both (default): *Incoming and outgoing traffic is blocked on a port configured for web authentication before authentication occurs.*

in: *Incoming traffic is blocked on a port configured for web authentication before authentication occurs. Outgoing traffic with unknown destination addresses is flooded on unauthenticated ports configured for web authentication.*

Prerequisites: *As implemented in 802.1X authentication, the disabling of incoming traffic and transmission of outgoing traffic on a web-authenticated egress port in an unauthenticated state (using the **aaa port-access controlled-directions in** command) is supported only if:*

- *The 802.1s Multiple Spanning Tree Protocol (MSTP) or 802.1w Rapid Spanning Tree Protocol (RSTP) is enabled on the switch. MSTP and RSTP improve resource utilization while maintaining a loop-free network.*

*The port is configured as an edge port in the network using the **spanning-tree edge-port** command.*

Syntax: aaa port-access <port-list> controlled-directions <both | in>
 — Continued —

Notes:

- *For information on how to configure the prerequisites for using the **aaa port-access controlled-directions in** command, see Chapter 4, “Multiple Instance Spanning-Tree Operation” in the Advanced Traffic Management Guide.*
- *To display the currently configured Controlled Directions value for web-authenticated ports, enter the **show port-access web-based config** command as shown in Figure 3-4.*
- *The **aaa port-access controlled-direction in** command allows Wake-on-LAN traffic to be transmitted on a web-authenticated egress port that has not yet transitioned to the authenticated state; the **controlled-direction both** setting prevents Wake-on-LAN traffic to be transmitted on a web-authenticated egress port until authentication occurs.*

The Wake-on-LAN feature is used by network administrators to remotely power on a sleeping workstation (for example, during early morning hours to perform routine maintenance operations, such as patch management and software updates)

- *Using the **aaa port-access controlled-directions in** command, you can enable the transmission of Wake-on-LAN traffic on unauthenticated egress ports that are configured for any of the following port-based security features:*
 - 802.1X authentication
 - MAC authentication
 - Web authentication

*Because a port can be configured for more than one type of authentication to protect the switch from unauthorized access, the last setting you configure with the **aaa port-access controlled-directions** command is applied to all authentication methods configured on the switch.*

For information about how to configure and use 802.1X authentication, refer to Chapter 10, “Configuring Port-Based and User-Based Access Control (802.1X)”.

- *When a web-authenticated port is configured with the **controlled-directions in** setting, eavesdrop prevention is not supported on the port.*

Syntax: [no] aaa port-access web-based <port-list>

*Enables web-based authentication on the specified ports. Use the **no** form of the command to disable web-based authentication on the specified ports.*

Syntax: aaa port-access web-based <port-list> [auth-vid <vid>]]

no aaa port-access web-based <port-list> [auth-vid]

*Specifies the VLAN to use for an authorized client. The Radius server can override the value (accept-response includes a vid). If **auth-vid** is 0, no VLAN changes occur unless the RADIUS server supplies one.*

*Use the **no** form of the command to set the **auth-vid** to 0. (Default: 0).*

Syntax: aaa port-access web-based [clear-statistics]

*Clears (resets to 0) all counters used to monitor the CEI, HTTP, Web-Auth control traffic generated in web authentication session. (To display Web-Auth traffic statistics, enter the **show port-access web-based statistics** command.)*

Syntax: aaa port-access web-based <port-list> [client-limit <1-32>]

Specifies the maximum number of authenticated clients to allow on the port. (Default: 1)

Note: *On switches where Web Auth and 802.1X can operate concurrently, this limit includes the total number of clients authenticated through both methods.*

Syntax: aaa port-access web-based [dhcp-addr <ip-address/mask>]

*Specifies the base address/mask for the temporary IP pool used by DHCP. The base address can be any valid ip address (not a multicast address). Valid mask range value is <255.255.240.0 - 255.255.255.0>.
(Default: 192.168.0.0/255.255.255.0)*

Syntax: aaa port-access web-based [dhcp-lease <5 - 25>]

*Specifies the lease length, in seconds, of the temporary IP address issued for Web Auth login purposes.
(Default: 10 seconds)*

Syntax: aaa port-access web-based [ewa-server <ipv4-addr | hostname>
 [<page-path>]]

Configures a connection with the web server at the specified IPv4 address (ipv4-addr) or host name (ipv4-addr) on which customized login web pages used for Web Authentication are stored. A maximum of 3 web servers may be configured on the switch.

The optional <page-path> parameter defines the directory path on the server where all customized login web pages (graphics, HTML frames, and HTML files) are stored. (Default: The default <page-path> value is "/" for root directory. If the web server is also used for other purposes, you may wish to group the HTML files in their own directory, for example in "/EWA/")

```
ProCurve Switch (config)# aaa port-access web-based 47 ewa-server 10.0.12.179
/EWA
ProCurve Switch (config)# aaa port-access web-based 47 ewa-server 10.0.12.180
/EWA
ProCurve Switch (config)#
```

Figure 6. Adding Web Servers with the aaa port-access web-based ewa-server Command

```
ProCurve Switch (config)# no aaa port-access web-based 47 ewa-server 10.0.12.181
ProCurve Switch (config)#
```

Figure 7. Removing a Web Server with the aaa port-access web-based ewa-server Command

Syntax: aaa port-access web-based <port-list> [logoff-period] <60-9999999>]

Specifies the period, in seconds, that the switch enforces for an implicit logoff. This parameter is equivalent to the MAC age interval in a traditional switch sense. If the switch does not see activity after a logoff-period interval, the client is returned to its pre-authentication state. (Default: 300 seconds)

Syntax: aaa port-access web-based <port-list> [max-requests <1-10>]

Specifies the number of authentication attempts that must time-out before authentication fails. (Default: 2)

Syntax: aaa port-access web-based <port-list> [max-retries <1-10>]

Specifies the number of the number of times a client can enter their user name and password before authentication fails. This allows the reentry of the user name and password if necessary. (Default: 3)

Syntax: aaa port-access web-based <port-list> [quiet-period <1 - 65535>]

Specifies the time period (in seconds) the switch uses before sending an authentication request for a client that failed authentication. (Default: 60 seconds)

Syntax: aaa port-access web-based <port-list> [reauth-period <0 - 9999999>]

Specifies the time period, in seconds, the switch enforces on a client to re-authenticate. When set to 0, reauthentication is disabled. (Default: 300 seconds)

Syntax: aaa port-access web-based <port-list> [reauthenticate]

Forces a reauthentication of all attached clients on the port.

Syntax: aaa port-access web-based <port-list> [redirect-url <url>]

no aaa port-access web-based <port-list> [redirect-url]

Specifies the URL that a user is redirected to after a successful login. Any valid, fully-formed URL may be used, for example, <http://welcome-server/welcome.htm> or <http://192.22.17.5>. ProCurve recommends that you provide a redirect URL when using Web Authentication.

Note: The **redirect-url** command accepts only the first 103 characters of the allowed 127 characters.

*Use the **no** form of the command to remove a specified redirect URL.*

(Default: There is no default URL. Browser behavior for authenticated clients may not be acceptable.)

Syntax: aaa port-access web-based [e] <port-list> [server-timeout <1 - 300>]

*Specifies the period, in seconds, the switch waits for a server response to an authentication request. Depending on the current **max-requests** value, the switch sends a new attempt or ends the authentication session. (Default: 30 seconds)*

Show Commands for Web Authentication

Command	Page
show port-access web-based [<i>port-list</i>]	3-28
show port-access web-based clients [<i>port-list</i>]	3-29
show port-access web-based clients < <i>port-list</i> > detailed	3-30
show port-access web-based config [<i>port-list</i>]	3-31
show port-access web-based config < <i>port-list</i> > detailed	3-32
show port-access web-based config [<i>port-list</i>] auth-server	3-33
show port-access web-based config [<i>port-list</i>] web-server	3-33

Syntax: show port-access web-based [*port-list*]

Displays the status of all ports or specified ports that are enabled for Web authentication. The information displayed for each port includes:

- *Number of authorized and unauthorized clients*
- *VLAN ID number of the untagged VLAN used. If the switch supports MAC-based (untagged) VLANs, **MACbased** is displayed to show that multiple untagged VLANs are configured for authentication sessions.*
- *If tagged VLANs (statically configured or RADIUS-assigned) are used (**Yes** or **No**)*
- *If client-specific per-port CoS (Class of Service) values are configured (**Yes** or **No**) or the numerical value of the CoS (802.1p priority) applied to all inbound traffic. For client-specific per-port CoS values, enter the **show port-access web-based clients detailed** command.*
- *If per-port rate-limiting for inbound traffic is applied (**Yes** or **No**) or the percentage value of the port's available bandwidth applied as a rate-limit value.*

Information on ports not enabled for Web Authentication is not displayed.

```
ProCurve (config)# show port-access web-based
```

Port Access Web-Based Status						
Port	Auth Clients	Unauth Clients	Untagged VLAN	Tagged VLANs	Port COS	Cntrl Dir
1	1	1	4006	Yes	70000000	
2	2	0	MACbased	No	Yes	
3	4	0	1	Yes	No	

Figure 4. Example of show port-access web-based Command Output

Syntax: show port-access web-based clients [*port-list*]

Displays the session status, name, and address for each web-authenticated client on the switch. The IP address displayed is taken from the DHCP binding table (learned through the DHCP Snooping feature).

*If DHCP snooping is not enabled on the switch, **n/a** (not available) is displayed for a client's IP address.*

*If a web-authenticated client uses an IPv6 address, **n/a - IPv6** is displayed.*

*If DHCP snooping is enabled but no MAC-to-IP address binding for a client is found in the DHCP binding table, **n/a - no info** is displayed.*

```
ProCurve (config)# show port-access web-based clients
```

Port Access Web-Based Client Status				
Port	Client Name	MAC Address	IP Address	Session Status
1	webuser1	0010b5-891a9e	192.192.192.192	Authenticated
1	webuser2	001560-b3ea48	n/a - no info	Authenticating
1	webuser3	000000-111111	n/a - IPv6	Authenticating
3	webuser4	000000-111112	n/a	Authenticating

Figure 5. Example of show port-access web-based clients Command Output

Syntax: show port-access web-based clients <port-list> detailed

Displays detailed information on the status of web-authenticated client sessions on specified switch ports.

```
ProCurve (config)# show port-access web-based clients 1 detailed

Port Access Web-Based Client Status Detailed

Client Base Details :
  Port           : 1
  Session Status : authenticated      Session Time(sec) : 6
  Username       : webuser1          MAC Address       : 0010b5-891a9e
  IP             : n/a

Access Policy Details :
  COS Map        : 11111111          In Limit %       : 98
  Untagged VLAN  : 4006              Out Limit %      : 100
  Tagged VLANs   : 1, 3, 5, 6, 334, 2566

RADIUS-ACL List :
  deny in udp from any to 10.2.8.233 CNT
    Hit Count: 0
  permit in udp from any to 10.2.8.233 CNT
    Hit Count: 0
  deny in tcp from any to 10.2.8.233 CNT
    Hit Count: 0
  permit in tcp from any to 10.2.8.233 CNT
    Hit Count: 0
  permit in tcp from any to 0.0.0.0/0 CNT
    Hit Count: 0
```

Figure 6. Example of show port-access web-based clients detailed Command Output

Syntax: show port-access web-based config [*port-list*]

Displays the currently configured Web Authentication settings for all switch ports or specified ports, including:

- *Temporary DHCP base address and mask*
- *Support for RADIUS-assigned dynamic VLANs (**Yes** or **No**)*
- *Controlled directions setting for transmitting Wake-on-LAN traffic on egress ports*
- *Authorized and unauthorized VLAN IDs*

*If the authorized or unauthorized VLAN ID value is **0**, the default VLAN ID is used unless overridden by a RADIUS-assigned value.*

ProCurve (config)# show port-access web-based config								
Port Access Web-Based Configuration								
DHCP Base Address : 192.168.0.0								
DHCP Subnet Mask : 255.255.255.0								
DHCP Lease Length : 10								
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No								
Port	Enabled	Client Limit	Client Moves	Logoff Period	Re-Auth Period	Unauth VLAN ID	Auth VLAN ID	Cntrl Dir
-----	-----	-----	-----	-----	-----	-----	-----	-----
1	Yes	1	No	300	0	0	0	both
2	Yes	1	No	300	0	0	0	in
...								

Figure 7. Example of show port-access web-based config Command Output

Syntax: show port-access web-based config <port-list> detailed

Displays more detailed information on the currently configured Web Authentication settings for specified ports.

```
ProCurve (config)# show port-access web-based config 1 detailed

Port Access Web-Based Detailed Configuration

Port           : 1           Web-based enabled : Yes
Client Limit   : 1           Client Moves      : No
Logoff Period  : 300         Re-Auth Period   : 0

Unauth VLAN ID : 0           Auth VLAN ID      : 0

Max Requests   : 3           Quiet Period      : 60
Server Timeout : 30

Max Retries     : 3           SSL Enabled       : No
Redirect URL    :
...
```

Figure 8. Example of show port-access web-based config detail Command Output

Syntax: show port-access web-based config [*port-list*] auth-server

Displays the currently configured Web Authentication settings for all switch ports or specified ports and includes RADIUS server-specific settings, such as:

- *Timeout waiting period*
- *Number of timeouts supported before authentication login fails*
- *Length of time (quiet period) supported between authentication login attempts*

```
ProCurve (config)# show port-access web-based config auth-server
```

Port Access Web-Based Configuration								
Port	Enabled	Client Limit	Client Moves	Logoff Period	Re-Auth Period	Max Req	Quiet Period	Server Timeout
1	Yes	1	No	300	0	3	60	30
2	No	1	No	300	0	3	60	30
...								

Figure 9. Example of show port-access web-based config auth-server Command Output

Syntax: show port-access web-based config [*port-list*] web-server

Displays the currently configured Web Authentication settings for all ports or specified ports, including web-specific settings for password retries, SSL login status, and a redirect URL, if specified.

Customizing Web Authentication HTML Files (Optional)

The Web Authentication process displays a series of web pages and status messages to the user during login. The web pages that are displayed can be:

- Generic, default pages generated directly by the switch software
- Customized pages hosted on a local web server.

By creating customized login web pages, you can improve the “look and feel” of the Web Authentication process to correspond more closely with your network and business needs. For example, you can:

- Identify the network that a client is trying to log into.
- Provide contact information if a client has difficulty connecting to the network.
- Incorporate CSS styles consistent with the appearance of your network.

Implementing Customized Web-Auth Pages

To implement “enhanced” Web Authentication pages, you need to:

- Configure and start a web server on your local network,
- Customize the HTML template files and make them accessible to the web server
- Configure the switch to display the customized files by using the **aaa port-access web-based ewa-server** command to specify the server’s IP address or host name and the path to the customized HTML files on the server. See “” on page 3-50 for syntax details.

Operating Notes and Guidelines

- Customized Web Authentication pages are configured per switch, so that each Web-Auth enabled port displays the same customized pages at client login.
- The customized web pages you create can be hosted on up to three web servers in your network. Implementing multiple web servers provides redundancy in case access to any of the other servers fail.

- To configure a web server on your network, follow the instructions in the documentation provided with the server.
- Before you enable custom Web Authentication pages, you should:
 - Determine the IP address or host name of the web server(s) that will host your custom pages.
 - Determine the path on the server(s) where the HTML files (including all graphics) used for the login pages are stored.
 - Configure and start the web server(s).
 - Create the customized web pages as described in “Customizing HTML Templates”, and store them in the document path on the designated servers.
 - Test that they are accessible at the designated URL(s).

Customizing HTML Templates

When you customize an HTML template, follow these guidelines:


- Do not change the name of any of the HTML files (**index.html**, **accept.html**, and so on).
- Some template pages use Embedded Switch Includes (ESIs) or Active Server Pages. These should not be modified when customizing HTML files. ESIs behave as follows:
 - i. A client’s web browser sends a request for an HTML file. The switch passes the request to a configured web server.
 - ii. The web server responds by sending a customized HTML page to the switch. Each ESI call in the HTML page is replaced with the value (in plain text) retrieved by the call.
 - iii. The switch sends the final version of the HTML page to the client’s web browser.
- Store all customized login web pages (including any graphics) that you create for client login on each web server at the path you will configure with the **aaa port-access web-based ewa-server** command.

Customizable HTML Templates

The sample HTML files described in the following sections are customizable templates. To help you create your own set HTML files, a set of the templates can be found on the download page for 'K' software.

File Name	Page
index.html	3-36
accept.html	3-38
authen.html	3-40
reject_unauthvlan.html	3-41
timeout.html	3-43
retry_login.html	3-44
sslredirect.html	3-46
rejectnovlan.html	3-48

User Login Page (index.html).



User Login

In order to access this network, you must first log in.

Username:

Password:

Figure 8. User Login Page

The **index.html** file is the first login page displayed, in which a client requesting access to the network enters a username and password. In the **index.html** Template file, you can customize any part of the source code except for the form that processes the username and password entered by a client.

```
<!--
ProCurve Web Authentication Template
index.html
-->
<html>

  <head>
    <title>User Login</title>
  </head>

  <body>
    <h1>User Login</h1>
    <p>In order to access this network, you must first log in.</p>

    <form action="/webauth/loginprocess" method="POST">
      <table>
        <tr>
          <td>Username: </td>
          <td><input name="user" type="text" /></td>
        </tr>
        <tr>
          <td>Password: </td>
          <td><input name="pass" type="password" /></td>
        </tr>
        <tr>
          <td></td>
          <td><input type="submit" value="Submit" /></td>
        </tr>
      </table>
    </form>

  </body>
</html>
```

Figure 9. HTML Code for User Login Page Template

Access Granted Page (accept.html).




Figure 9-10. Access Granted Page

The **accept.html** file is the web page used to confirm a valid client login. This web page is displayed after a valid username and password are entered and accepted.

The client device is then granted access to the network. To configure the VLAN used by authorized clients, specify a VLAN ID with the **aaa port-access web-based auth-vid** command parameter when you enable Web Authentication.

The **accept.html** file contains the following ESIs, which should not be modified:

- The WAUTHREDIRECTTIMEGET ESI inserts the value for the waiting time used by the switch to redirect an authenticated client while the client renews its IP address and gains access to the network.
- The WAUTHREDIRECTURLGET ESI inserts the URL configured with the **redirect-url** parameter to redirect a client login or the first web page requested by the client.

```
<!--
ProCurve Web Authentication Template
accept.html
-->
<html>
  <head>
    <title>Access Granted</title>

    <!-- The following line is required to automatically redirect -->
    <meta http-equiv="refresh" content="<!-- ESI(WAUTHREDIRECTTIMEGET, 1) ->;URL=<
ESI(WAUTHREDIRECTURLGET, 1) ->" />
  </head>

  <body>
    <h1>Access Granted</h1>
    <!--
      The ESI tag below will be replaced with the time in seconds until
      the page redirects.
    -->
    <p>You have been authenticated. Please wait <!-- ESI(WAUTHREDIRECTTIMEGET, 1
-> seconds while network connection refreshes itself.</p>
  </body>
</html>
```

Figure 11. HTML Code for Access Granted Page Template

Authenticating Page (authen.html).




Figure 12. Authenticating Page

The **authen.html** file is the web page used to process a client login and is refreshed while user credentials are checked and verified.

```
<!--
ProCurve Web Authentication Template
authen.html
-->
<html>

  <head>
    <title>Authenticating</title>

    <!-- The following line is always required -->
    <meta http-equiv="refresh" content="2;URL=/webauth/statusprocess">
  </head>

  <body>
    <h1>Authenticating...</h1>
    <p>Please wait while your credentials are verified.</p>
  </body>

</html>
```

Figure 13. HTML Code for Authenticating Page Template

Invalid Credentials Page (reject_unauthvlan.html).

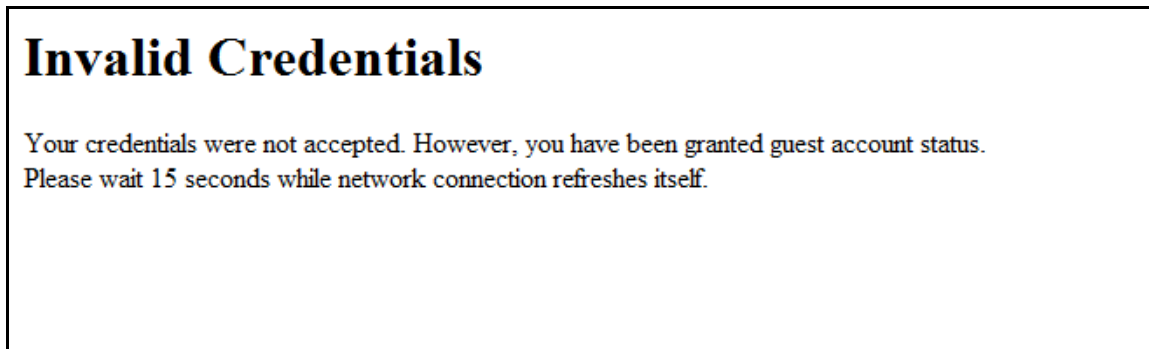


Figure 10. Invalid Credentials Page

The **reject_unauthvlan.html** file is the web page used to display login failures in which an unauthenticated client is assigned to the VLAN configured for unauthorized client sessions. You can configure the VLAN used by unauthorized clients with the **aaa port-access web-based unauth-vid** command when you enable Web Authentication.

The **WAUTHREDIRECTTIMEGET** ESI inserts the value for the waiting time used by the switch to redirect an unauthenticated client while the client renews its IP address and gains access to the VLAN for unauthorized clients. This ESI should not be modified.

```
<!--
ProCurve Web Authentication Template
reject_unauthvlan.html
-->
<html>
  <head>
    <title>Invalid Credentials</title>

    <!-- The following line is required to automatically redirect -->
    <meta http-equiv="refresh" content="<!-- ESI(WAUTHREDIRECTTIMEGET, 1) ->;URL=<
ESI(WAUTHREDIRECTURLGET, 1) ->" />
  </head>

  <body>
    <h1>Invalid Credentials</h1>
    <p>Your credentials were not accepted. However, you have been granted guest
account status. Please wait <!-- ESI(WAUTHREDIRECTTIMEGET, 1) -> seconds while network
connection refreshes itself.</p>
  </body>
</html>
```

Figure 14. HTML Code for Invalid Credentials Page Template

Timeout Page (timeout.html).




Figure 15. Timeout Page

The **timeout.html** file is the web page used to return an error message if the RADIUS server is not reachable. You can configure the time period (in seconds) that the switch waits for a response from the RADIUS server used to verify client credentials with the **aaa port-access web-based server-timeout** command when you enable Web Authentication.

```
<!--
ProCurve Web Authentication Template
timeout.html
-->
<html>

  <head>
    <title>Timeout</title>
  </head>

  <body>
    <h1>Timeout</h1>
    <p>Your credentials could not be verified with authentication server.
Please retry later.</p>
  </body>

</html>
```

Figure 16. HTML Code for Timeout Page Template

Retry Login Page (retry_login.html).




Figure 17. Retry Login Page

The **retry_login.html** file is the web page displayed to a client that has entered an invalid username and/or password, and is given another opportunity to log in.

The **WAUTHRETRIESLEFTGET** ESI displays the number of login retries that remain for a client that entered invalid login credentials. You can configure the number of times that a client can enter their user name and password before authentication fails with the **aaa port-access web-based max-retries** commands when you enable Web Authentication. This ESI should not be modified.

```
<!--
ProCurve Web Authentication Template
retry_login.html
-->
<html>

  <head>
    <title>Invalid Credentials</title>

    <!--
      The following line is required to automatically redirect
      the user back to the login page.
    -->
    <meta http-equiv="refresh" content="5;URL=/EWA/index.html">
  </head>

  <body>
    <h1>Invalid Credentials</h1>
    <p>Your credentials were not accepted. You have <!-- ESI(WAUTHRETRIESLEFTGET,1
-> retries left. Please try again.</p>
  </body>

</html>
```

Figure 18. HTML Code for Retry Login Page Template

SSL Redirect Page (sslredirect.html).

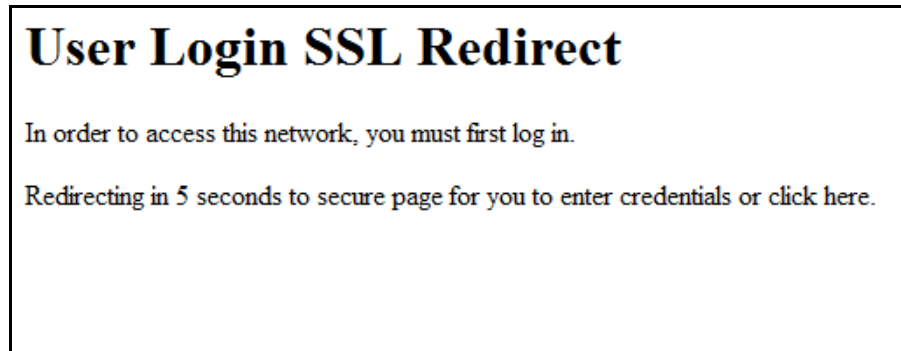


Figure 19. SSL Redirect Page

The **sslredirect** file is the web page displayed when a client is redirected to an SSL server to enter credentials for Web Authentication. If you have enabled SSL on the switch, you can enable secure SSL-based Web Authentication by entering the **aaa port-access web-based ssl-login** command when you enable Web Authentication.

The WAUTHSSLSRVGET ESI inserts the URL that redirects a client to an SSL-enabled port on a server to verify the client's username and password. This ESI should not be modified.

```
<!--  
ProCurve Web Authentication Template  
sslredirect.html  
-->  
<html>  
  
  <head>  
    <title>User Login SSL Redirect</title>  
  
    <meta http-equiv="refresh" content="5;URL=https://<!-- ESI(WAUTHSSLSRVGET,1  
->/EWA/index.html">  
  </head>  
  
  <body>  
    <h1>User Login SSL Redirect</h1>  
    <p>In order to access this network, you must first log in.</p>  
    <p>Redirecting in 5 seconds to secure page for you to enter credentials or <  
href="https://<!-- ESI(WAUTHSSLSRVGET,1) ->/EWA/index.html">click here</a>.</p>  
  </body>  
  
</html>
```

Figure 20. HTML Code for SSL Redirect Page Template

Access Denied Page (reject_novlan.html).

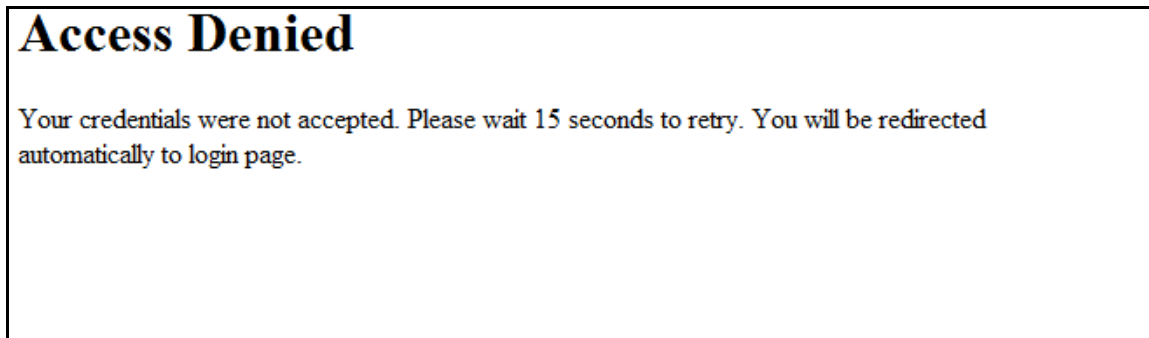


Figure 11. Access Denied Page

The **reject_novlan** file is the web page displayed after a client login fails and no VLAN is configured for unauthorized clients.

The WAUTHQUIETTIMEGET ESI inserts the time period used to block an unauthorized client from attempting another login. To specify the time period before a new authentication request can be received by the switch, configure a value for the **aaa port-access web-based quiet-period** command when you enable Web Authentication. This ESI should not be modified.

```
<!--
ProCurve Web Authentication Template
reject_novlan.html
-->
<html>

  <head>
    <title>Access Denied</title>

    <!--
      The line below is required to automatically redirect the user
      back to the login page.
    -->
    <meta http-equiv="refresh" content="<!-- ESI(WAUTHQUIETTIMEGET,1) ->;URL=/EW
index.html">
  </head>

  <body>
    <h1>Access Denied</h1>
    <p>Your credentials were not accepted. Please wait <!-- ESI(WAUTHQUIETTIMEGET
1) -> seconds to retry. You will be redirected automatically to login page.</p>
  </body>

</html>
```

Figure 21. HTML Code for Access Denied Page Template

Configuring MAC Authentication on the Switch

Overview

1. If you have not already done so, configure a local username and password pair on the switch.
2. If you plan to use multiple VLANs with MAC Authentication, ensure that these VLANs are configured on the switch and that the appropriate port assignments have been made.
3. Use the **ping** command in the switch console interface to ensure that the switch can communicate with the RADIUS server you have configured to support MAC-Auth on the switch.
4. Configure the switch with the correct IP address and encryption key to access the RADIUS server.
5. Configure the switch for MAC-Auth:
 - a. Configure MAC Authentication on the switch ports you want to use.
6. Test both the authorized and unauthorized access to your system to ensure that MAC Authentication works properly on the ports you have configured for port-access.

Configuration Commands for MAC Authentication

Command	Page
Configuration Level	
aaa port-access mac-based addr-format	3-51
[no] aaa port-access mac-based [e] <port-list>	3-52
[addr-limit]	3-52
[addr-moves]	3-52
[auth-vid]	3-52
[logoff-period]	3-53
[max-requests]	3-53
[quiet-period]	3-53
[reauth-period]	3-53
[reauthenticate]	3-53
[server-timeout]	3-53
[unauth-vid]	3-54

Syntax: aaa port-access mac-based addr-format <no-delimiter | single-dash | multi-dash | multi-colon | no-delimiter-uppercase | single-dash-uppercase | multi-dash-uppercase | multi-colon-uppercase>

Specifies the MAC address format to be used in the RADIUS request message. This format must match the format used to store the MAC addresses in the RADIUS server. (Default: no-delimiter)

no-delimiter — specifies an aabbccddeeff format.

single-dash — specifies an aabbcc-ddeeff format.

multi-dash — specifies an aa-bb-cc-dd-ee-ff format.

multi-colon — specifies an aa:bb:cc:dd:ee:ff format.

no-delimiter-uppercase — specifies an AABBCDDEEFF format.

single-dash-uppercase — specifies an AABCC-DDEEFF format

multi-dash-uppercase — specifies an AA-BB-CC-DD-EE-FF format

multi-colon-uppercase — specifies an AA:BB:CC:DD:EE:FF format.

Syntax: [no] aaa port-access mac-based < port-list >

*Enables MAC-based authentication on the specified ports. Use the **no** form of the command to disable MAC-based authentication on the specified ports.*

Syntax: aaa port-access mac-based [e] < port-list > [addr-limit <1-32>]

Specifies the maximum number of authenticated MACs to allow on the port. (Default: 1)

Note: On switches where MAC Auth and 802.1X can operate concurrently, this limit includes the total number of clients authenticated through both methods.

Syntax: [no] aaa port-access mac-based [e] < port-list > [addr-moves]

*Allows client moves between the specified ports under MAC Auth control. When enabled, the switch allows addresses to move without requiring a re-authentication. When disabled, the switch does not allow moves and when one does occur, the user will be forced to re-authenticate. At least two ports (from port(s) and to port(s)) must be specified. Use the **no** form of the command to disable MAC address moves between ports under MAC Auth control. (Default: disabled – no moves allowed)*

Syntax: aaa port-access mac-based [e] < port-list > [auth-vid <vid>]

no aaa port-access mac-based [e] < port-list > [auth-vid]

*Specifies the VLAN to use for an authorized client. The Radius server can override the value (accept-response includes a vid). If **auth-vid** is 0, no VLAN changes occur unless the RADIUS server supplies one. Use the **no** form of the command to set the **auth-vid** to 0. (Default: 0).*

Syntax: aaa port-access mac-based [e] <port-list>
[logoff-period] <60-9999999>]

Specifies the period, in seconds, that the switch enforces for an implicit logoff. This parameter is equivalent to the MAC age interval in a traditional switch sense. If the switch does not see activity after a logoff-period interval, the client is returned to its pre-authentication state. (Default: 300 seconds)

Syntax: aaa port-access mac-based [e] <port-list> [max-requests <1-10>]

Specifies the number of authentication attempts that must time-out before authentication fails. (Default: 2)

Syntax: aaa port-access mac-based [e] <port-list> [quiet-period <1 - 65535>]

Specifies the time period (in seconds) that the switch waits before processing an authentication request from a MAC address that failed authentication. (Default: 60 seconds)

Syntax: aaa port-access mac-based [e] <port-list> [reauth-period <0 - 9999999>]

Specifies the time period (in seconds) that the switch enforces on a client to re-authenticate. The client remains authenticated while the reauthentication occurs.

When set to **0**, reauthentication is disabled. (Default: 300 seconds)

Syntax: aaa port-access mac-based [e] < *port-list* > [reauthenticate]

Forces a reauthentication of all attached clients on the port.

Syntax: aaa port-access mac-based [e] <port-list> [server-timeout <1 - 300>]

*Specifies the period, in seconds, the switch waits for a server response to an authentication request. Depending on the current **max-requests** value, the switch sends a new attempt or ends the authentication session. (Default: 30seconds)*

Syntax: aaa port-access mac-based [e] <port-list> [unauth-vid <vid>]
 no aaa port-access mac-based [e] <port-list> [unauth-vid]

*Specifies the VLAN to use for a client that fails authentication. If **unauth-vid** is **0**, no VLAN changes occur. Use the **no** form of the command to set the **unauth-vid** to **0**. (Default: 0)*

Show Commands for MAC-Based Authentication

Command	Page
show port-access mac-based [port-list]	3-55
show port-access mac-based clients [port-list]	3-55
show port-access mac-based clients <port-list> detailed	3-56
show port-access mac-based config [port-list]	3-57
show port-access mac-based config <port-list> detailed	3-58
show port-access mac-based config [port-list] auth-server	3-59

Syntax: show port-access mac-based [port-list]

Displays the status of all ports or specified ports that are enabled for MAC authentication. The information displayed for each port includes:

- *Number of authorized and unauthorized clients*
- *VLAN ID number of the untagged VLAN used. If the switch supports MAC-based (untagged) VLANs, **MACbased** is displayed to show that multiple untagged VLANs are configured for authentication sessions.*
- *If tagged VLANs (statically configured or RADIUS-assigned) are used (**Yes** or **No**)*
- *If client-specific per-port CoS (Class of Service) values are configured (**Yes** or **No**) or the numerical value of the CoS (802.1p priority) applied to all inbound traffic. For client-specific per-port CoS values, enter the **show port-access web-based clients detailed** command.*
- *If per-port rate-limiting for inbound traffic is applied (**Yes** or **No**) or the percentage value of the port's available bandwidth applied as a rate-limit value.*

Information on ports not enabled for MAC authentication is not displayed.

```
ProCurve (config)# show port-access mac-based
```

Port Access MAC-Based Status						
Port	Auth Clients	Unauth Clients	Untagged VLAN	Tagged VLANs	Port COS	Cntrl Dir
1	1	1	2003	Yes	70000000	
2	2	0	MACbased	No	Yes	
3	4	0	1	Yes	No	

Figure 3-22. Example of show port-access mac-based Command Output

Syntax: show port-access mac-based clients [port-list]

Displays the session status, name, and address for each MAC-authenticated client on the switch. The IP address displayed is taken from the DHCP binding table (learned through the DHCP Snooping feature).

*If DHCP snooping is not enabled on the switch, **n/a** (not available) is displayed for a client's IP address.*

*If a MAC-authenticated client uses an IPv6 address, **n/a - IPv6** is displayed.*

*If DHCP snooping is enabled but no MAC-to-IP address binding for a client is found in the DHCP binding table, **n/a - no info** is displayed.*

```
ProCurve (config)# show port-access mac-based clients
```

Port Access MAC-Based Client Status			
Port	MAC Address	IP Address	Session Status
1	001321-eb8063	2001:fece:ba23:cd1f:dcb1:1010:9234:4088	unauthenticated
1	000000-111111	192.192.192.192	authenticated
2	000000-111111	n/a	authenticating

Figure 4. Example of show port-access mac-based clients Command Output

Syntax: show port-access mac-based clients <port-list> detailed
Displays detailed information on the status of MAC-authenticated client sessions on specified ports.

```
ProCurve (config)# show port-access mac-based clients 1 detailed

Port Access MAC-Based Client Status Detailed

Client Base Details :
  Port           : 1
  Session Status : authenticated      Session Time(sec) : 6
  Username       : client1            MAC Address       : 0010b5-891a9e
  IP             : n/a

Access Policy Details :
  COS Map        : 12345678           In Limit %       : 98
  Untagged VLAN  : 4006               Out Limit %      : 100
  Tagged VLANs   : 1, 3, 5, 6, 334, 4001

RADIUS-ACL List :
  deny in udp from any to 10.2.8.233 CNT
    Hit Count: 0
  permit in udp from any to 10.2.8.233 CNT
    Hit Count: 0
  deny in tcp from any to 10.2.8.233 CNT
    Hit Count: 0
  permit in tcp from any to 10.2.8.233 CNT
    Hit Count: 0
  permit in tcp from any to 0.0.0.0/0 CNT
    Hit Count: 0
```

Figure 5. Example of show port-access mac-based clients detail Command Output

Syntax: show port-access mac-based config [port-list]

Displays the currently configured MAC Authentication settings for all switch ports or specified ports, including:

- *MAC address format*
- *Support for RADIUS-assigned dynamic VLANs (**Yes** or **No**)*
- *Controlled directions setting for transmitting Wake-on-LAN traffic on egress ports*
- *Authorized and unauthorized VLAN IDs*

*If the authorized or unauthorized VLAN ID value is **0**, the default VLAN ID is used unless overridden by a RADIUS-assigned value.*

```
ProCurve (config)# show port-access mac-based config

Port Access MAC-Based Configuration

MAC Address Format : no-delimiter
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

  Port  Enabled  Client Limit  Client Moves  Logoff Period  Re-Auth Period  Unauth VLAN ID  Auth VLAN ID  Cntrl Dir
  ----  -
1      No      1           No            300            0              0              0             both
2      Yes     1           No            300            0              0              0             in
...
```

Figure 6. Example of show port-access mac-based config Command Output

Syntax: show port-access mac-based config <port-list> detailed

Displays more detailed information on the currently configured MAC Authentication settings for specified ports.

```
ProCurve (config)# show port-access mac-based config 1 detailed

Port Access MAC-Based Detailed Configuration

Port           : 1           Web-based enabled : Yes
Client Limit   : 1           Client Moves      : No
Logoff Period  : 300         Re-Auth Period   : 0

Unauth VLAN ID : 0           Auth VLAN ID      : 0

Max Requests   : 3           Quiet Period      : 60
Server Timeout : 30
```

Figure 7. Example of show port-access mac-based config detail Command Output

Syntax: show port-access mac-based config [*port-list*] auth-server

Displays the currently configured Web Authentication settings for all switch ports or specified ports and includes RADIUS server-specific settings, such as:

- *Timeout waiting period*
- *Number of timeouts supported before authentication login fails*
- *Length of time (quiet period) supported between authentication login attempts*

```
ProCurve (config)# show port-access mac-based config auth-server
```

Port Access MAC-Based Configuration								
Port	Enabled	Client Limit	Client Moves	Logoff Period	Re-Auth Period	Max Req	Quiet Period	Server Timeout
1	No	1	No	300	0	3	60	30
2	No	1	No	300	0	3	60	30
3	Yes	1	No	300	0	3	60	30
...								

Figure 8. Example of show port-access mac-based config auth-server Command Output

Client Status

The table below shows the possible client status information that may be reported by a Web-based or MAC-based **'show... clients'** command.

Reported Status	Available Network Connection	Possible Explanations
authenticated	Authorized VLAN	Client authenticated. Remains connected until logoff-period or reauth-period expires.
authenticating	Switch only	Pending RADIUS request.
rejected-no vlan	No network access	<ol style="list-style-type: none">1. Invalid credentials supplied.2. RADIUS Server difficulties. See log file.3. If unauth-vid is specified it cannot be successfully applied to the port. An authorized client on the port has precedence.
rejected-unauth vlan	Unauthorized VLAN only	<ol style="list-style-type: none">1. Invalid credentials supplied.2. RADIUS Server difficulties. See log file.
timed out-no vlan	No network access	RADIUS request timed out. If unauth-vid is specified it cannot be successfully applied to the port. An authorized client on the port has precedence. Credentials resubmitted after quiet-period expires.
timed out-unauth vlan	Unauthorized VLAN only	RADIUS request timed out. After the quiet-period expires credentials are resubmitted when client generates traffic.
unauthenticated	Switch only	Waiting for user credentials.

TACACS+ Authentication

Contents

Overview	4-2
Terminology Used in TACACS Applications:	4-3
General System Requirements	4-5
General Authentication Setup Procedure	4-5
Configuring TACACS+ on the Switch	4-8
Before You Begin	4-8
CLI Commands Described in this Section	4-9
Viewing the Switch's Current Authentication Configuration	4-9
Viewing the Switch's Current TACACS+ Server Contact Configuration	4-10
Configuring the Switch's Authentication Methods	4-11
Using the Privilege-Mode Option for Login	4-11
Authentication Parameters	4-12
Configuring the TACACS+ Server for Single Login	4-13
Configuring the Switch's TACACS+ Server Access	4-18
How Authentication Operates	4-25
General Authentication Process Using a TACACS+ Server	4-25
Local Authentication Process	4-27
Using the Encryption Key	4-28
General Operation	4-28
Encryption Options in the Switch	4-28
Controlling Web Browser Interface Access When Using TACACS+ Authentication	4-29
Messages Related to TACACS+ Operation	4-30
Operating Notes	4-30

Overview

Feature	Default	Menu	CLI	Web
view the switch's authentication configuration	n/a	—	page 4-9	—
view the switch's TACACS+ server contact configuration	n/a	—	page 4-10	—
configure the switch's authentication methods	disabled	—	page 4-11	—
configure the switch to contact TACACS+ server(s)	disabled	—	page 4-18	—

TACACS+ authentication enables you to use a central server to allow or deny access to the switches covered in this guide (and other TACACS-aware devices) in your network. This means that you can use a central database to create multiple unique username/password sets with associated privilege levels for use by individuals who have reason to access the switch from either the switch's console port (local access) or Telnet (remote access).

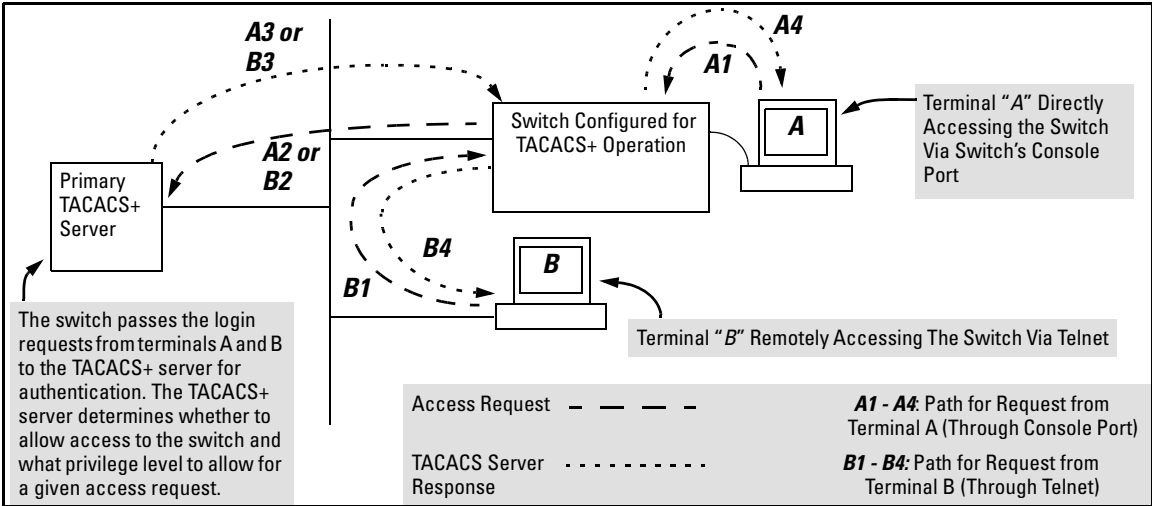


Figure 4-1. Example of TACACS+ Operation

TACACS+ in the switches covered in this guide manages authentication of logon attempts through either the Console port or Telnet. TACACS+ uses an authentication hierarchy consisting of (1) remote passwords assigned in a TACACS+ server and (2) local passwords configured on the switch. That is, with TACACS+ configured, the switch first tries to contact a designated

TACACS+ server for authentication services. If the switch fails to connect to any TACACS+ server, it defaults to its own locally assigned passwords for authentication control if it has been configured to do so. For both Console and Telnet access you can configure a login (read-only) and an enable (read/write) privilege level access.

TACACS+ does not affect web browser interface access. See “Controlling Web Browser Interface Access” on page 4-28.

Terminology Used in TACACS Applications:

- **NAS (Network Access Server):** This is an industry term for a TACACS-aware device that communicates with a TACACS server for authentication services. Some other terms you may see in literature describing TACACS operation are *communication server*, *remote access server*, or *terminal server*. These terms apply to a switch when TACACS+ is enabled on the switch (that is, when the switch is TACACS-aware).
- **TACACS+ Server:** The server or management station configured as an access control server for TACACS-enabled devices. To use TACACS+ with a switch covered in this guide and any other TACACS-capable devices in your network, you must purchase, install, and configure a TACACS+ server application on a networked server or management station in the network. The TACACS+ server application you install will provide various options for access control and access notifications. For more on the TACACS+ services available to you, see the documentation provided with the TACACS+ server application you will use.
- **Authentication:** The process for granting user access to a device through entry of a user name and password and comparison of this username/password pair with previously stored username/password data. Authentication also grants levels of access, depending on the privileges assigned to a user name and password pair by a system administrator.
 - **Local Authentication:** This method uses username/password pairs configured locally on the switch; one pair each for manager-level and operator-level access to the switch. You can assign local usernames and passwords through the CLI or web browser inter-

face. (Using the menu interface you can assign a local password, but not a username.) Because this method assigns passwords to the switch instead of to individuals who access the switch, you must distribute the password information on each switch to everyone who needs to access the switch, and you must configure and manage password protection on a per-switch basis. (For more on local authentication, refer to chapter 2, “Configuring Username and Password Security”.)

- **TACACS+ Authentication:** This method enables you to use a TACACS+ server in your network to assign a unique password, user name, and privilege level to each individual or group who needs access to one or more switches or other TACACS-aware devices. This allows you to administer primary authentication from a central server, and to do so with more options than you have when using only local authentication. (You will still need to use local authentication as a backup if your TACACS+ servers become unavailable.) This means, for example, that you can use a central TACACS+ server to grant, change, or deny access to a specific individual on a specific switch instead of having to change local user name and password assignments on the switch itself, and then have to notify other users of the change.

General System Requirements

To use TACACS+ authentication, you need the following:

- A TACACS+ server application installed and configured on one or more servers or management stations in your network. (There are several TACACS+ software packages available.)
- A switch configured for TACACS+ authentication, with access to one or more TACACS+ servers.

Notes

The effectiveness of TACACS+ security depends on correctly using your TACACS+ server application. For this reason, ProCurve recommends that you thoroughly test all TACACS+ configurations used in your network.

TACACS-aware ProCurve switches include the capability of configuring multiple backup TACACS+ servers. ProCurve recommends that you use a TACACS+ server application that supports a redundant backup installation. This allows you to configure the switch to use a backup TACACS+ server if it loses access to the first-choice TACACS+ server.

TACACS+ does not affect web browser interface access. Refer to “Controlling Web Browser Interface Access When Using TACACS+ Authentication” on page 4-28.

General Authentication Setup Procedure

It is important to test the TACACS+ service before fully implementing it. Depending on the process and parameter settings you use to set up and test TACACS+ authentication in your network, you could accidentally lock all users, including yourself, out of access to a switch. While recovery is simple, it may pose an inconvenience that can be avoided. To prevent an unintentional lockout on the switch, use a procedure that configures and tests TACACS+ protection for one access type (for example, Telnet access), while keeping the

other access type (console, in this case) open in case the Telnet access fails due to a configuration problem. The following procedure outlines a general setup procedure.

Note

If a complete access lockout occurs on the switch as a result of a TACACS+ configuration, see “Troubleshooting TACACS+ Operation” in the Troubleshooting chapter of the *Management and Configuration Guide* for your switch.

1. Familiarize yourself with the requirements for configuring your TACACS+ server application to respond to requests from the switch. (Refer to the documentation provided with the TACACS+ server software.) This includes knowing whether you need to configure an encryption key. (See “Using the Encryption Key” on page 4-27.)
2. Determine the following:
 - The IP address(es) of the TACACS+ server(s) you want the switch to use for authentication. If you will use more than one server, determine which server is your first-choice for authentication services.
 - The encryption key, if any, for allowing the switch to communicate with the server. You can use either a global key or a server-specific key, depending on the encryption configuration in the TACACS+ server(s).
 - The number of log-in attempts you will allow before closing a log-in session. (Default: 3)
 - The period you want the switch to wait for a reply to an authentication request before trying another server.
 - The username/password pairs you want the TACACS+ server to use for controlling access to the switch.
 - The privilege level you want for each username/password pair administered by the TACACS+ server for controlling access to the switch.
 - The username/password pairs you want to use for local authentication (one pair each for Operator and Manager levels).
3. Plan and enter the TACACS+ server configuration needed to support TACACS+ operation for Telnet access (login and enable) to the switch. This includes the username/password sets for logging in at the Operator (read-only) privilege level and the sets for logging in at the Manager (read/write) privilege level.

**Note on
Privilege Levels**

When a TACACS+ server authenticates an access request from a switch, it includes a privilege level code for the switch to use in determining which privilege level to grant to the terminal requesting access. The switch interprets a privilege level code of “15” as authorization for the Manager (read/write) privilege level access. Privilege level codes of 14 and lower result in Operator (read-only) access. Thus, when configuring the TACACS+ server response to a request that includes a username/password pair that should have Manager privileges, you must use a privilege level of 15. For more on this topic, refer to the documentation you received with your TACACS+ server application.

If you are a first-time user of the TACACS+ service, ProCurve recommends that you configure only the minimum feature set required by the TACACS+ application to provide service in your network environment. After you have success with the minimum feature set, you may then want to try additional features that the application offers.

4. Ensure that the switch has the correct local username and password for Manager access. (If the switch cannot find any designated TACACS+ servers, the local manager and operator username/password pairs are always used as the secondary access control method.)

Caution

You should ensure that the switch has a local Manager password. Otherwise, if authentication through a TACACS+ server fails for any reason, then unauthorized access will be available through the console port or Telnet.

5. Using a terminal device connected to the switch’s console port, configure the switch for TACACS+ authentication *only* for **telnet login** access and **telnet enable** access. At this stage, do not configure TACACS+ authentication for console access to the switch, as you may need to use the console for access if the configuration for the Telnet method needs debugging.
6. Ensure that the switch is configured to operate on your network and can communicate with your first-choice TACACS+ server. (At a minimum, this requires IP addressing and a successful **ping** test from the switch to the server.)
7. On a remote terminal device, use Telnet to attempt to access the switch. If the attempt fails, use the console access to check the TACACS+ configuration on the switch. If you make changes in the switch configuration, check Telnet access again. If Telnet access still fails, check the

configuration in your TACACS+ server application for mis-configurations or missing data that could affect the server's interoperability with the switch.

8. After your testing shows that Telnet access using the TACACS+ server is working properly, configure your TACACS+ server application for console access. Then test the console access. If access problems occur, check for and correct any problems in the switch configuration, and then test console access again. If problems persist, check your TACACS+ server application for mis-configurations or missing data that could affect the console access.
9. When you are confident that TACACS+ access through both Telnet and the switch's console operates properly, use the **write memory** command to save the switch's running-config file to flash.

Configuring TACACS+ on the Switch

Before You Begin

If you are new to TACACS+ authentication, ProCurve recommends that you read the "General Authentication Setup Procedure" on page 4-5 and configure your TACACS+ server(s) before configuring authentication on the switch.

The switch offers three command areas for TACACS+ operation:

- **show authentication** and **show tacacs**: Displays the switch's TACACS+ configuration and status.
- **aaa authentication**: A command for configuring the switch's authentication methods
- **tacacs-server**: A command for configuring the switch's contact with TACACS+ servers

CLI Commands Described in this Section

Command	Page
show authentication	4-9
show tacacs	4-10
aaa authentication	4-11 through 4-17
console	
Telnet	
num-attempts <1-10 >	
tacacs-server	4-18
host < ip-addr >	4-18
key	4-22
timeout < 1-255 >	4-23

Viewing the Switch's Current Authentication Configuration

This command lists the number of login attempts the switch allows in a single login session, and the primary/secondary access methods configured for each type of access.

Syntax: show authentication

This example shows the default authentication configuration.

```
ProCurve> show authentication
Status and Counters - Authentication Information
Login Attempts : 3

      Login   Login   Enable  Enable
Access Task Primary Secondary Primary Secondary
-----
(Console) Local  None    Local  None
(Telnet)  Local  None    Local  None
```

Configuration for login and enable access to the switch through the switch console port.

Configuration for login and enable access to the switch through Telnet.


Figure 4-2. Example Listing of the Switch's Authentication Configuration

Viewing the Switch's Current TACACS+ Server Contact Configuration

This command lists the timeout period, encryption key, and the IP addresses of the first-choice and backup TACACS+ servers the switch can contact.

Syntax: show tacacs

For example, if the switch was configured for a first-choice and two backup TACACS+ server addresses, the default timeout period, and **paris-1** for a (global) encryption key, **show tacacs** would produce a listing similar to the following:



ProCurve# show tacacs									
Status and Counters - TACACS Information									
Timeout : 5									
Encryption Key : paris-1									
First-Choice TACACS+ Server	Server	IP Addr	Opens	Closes	Aborts	Errors	Pkts Rx	Pkts Tx	
	-----	-----	-----	-----	-----	-----	-----	-----	
	10.30.248.100		0	0	0	0	0	0	
Second-Choice TACACS+ Server	10.30.248.156		0	0	0	0	0	0	
Third-Choice TACACS+ Server	10.30.248.105		0	0	0	0	0	0	

Figure 4-3. Example of the Switch's TACACS+ Configuration Listing

Configuring the Switch's Authentication Methods

The **aaa authentication** command configures access control for the following access methods:

- Console
- Telnet
- SSH
- Web
- Port-access (802.1X)

However, TACACS+ authentication is only used with the console, Telnet, or SSH access methods. The command specifies whether to use a TACACS+ server or the switch's local authentication, or (for some secondary scenarios) no authentication (meaning that if the primary method fails, authentication is denied). The command also reconfigures the number of access attempts to allow in a session if the first attempt uses an incorrect username/password pair.

Using the Privilege-Mode Option for Login

When using TACACS+ to control user access to the switch, you must first login with your username at the Operator privilege level using the password for Operator privileges, and then login again with the same username but using the Manager password to obtain Manager privileges. You can avoid this double login process by entering the **privilege-mode** option with the **aaa authentication login** command to enable TACACS+ for a single login. The switch authenticates your username/password, then requests the privilege level (Operator or Manager) that was configured on the TACACS+ server for this username/password. The TACACS+ server returns the allowed privilege level to the switch. You are placed directly into Operator or Manager mode, depending on your privilege level.

```
ProCurve(config) aaa authentication login privilege-mode
```

The **no** version of the above command disables TACACS+ single login capability.

Syntax: aaa authentication

< console | telnet | ssh | web | port-access >

Selects the access method for configuration.

< enable >

The server grants privileges at the Manager privilege level.

< login [privilege-mode] >

*The server grants privileges at the Operator privilege level. If the **privilege-mode** option is entered, TACACS+ is enabled for a single login. The authorized privilege level (Operator or Manager) is returned to the switch by the TACACS+ server.*

Default: Single login disabled.

< local | tacacs | radius >

Selects the type of security access:

local — Authenticates with the Manager and Operator password you configure in the switch.

tacacs — Authenticates with a password and other data configured on a TACACS+ server.

radius — Authenticates with a password and other data configured on a RADIUS server.

[< local | none >]

If the primary authentication method fails, determines whether to use the local password as a secondary method or to disallow access.

aaa authentication num-attempts < 1-10 >

Specifies the maximum number of login attempts allowed in the current session. Default: 3

Authentication Parameters

Table 4-1. AAA Authentication Parameters


Name	Default	Range	Function
console, Telnet, SSH, web or port-access	n/a	n/a	Specifies the access method used when authenticating. TACACS+ authentication only uses the console, Telnet or SSH access methods.

Name	Default	Range	Function
enable	n/a	n/a	Specifies the Manager (read/write) privilege level for the access method being configured.
login <privilege-mode>	privilege-mode disabled	n/a	<p>login: Specifies the Operator (read-only) privilege level for the access method being configured.</p> <p>The privilege-mode option enables TACACS+ for a single login. The authorized privilege level (Operator or Manager) is returned to the switch by the TACACS+ server.</p>
local - or - tacacs	local	n/a	<p>Specifies the primary method of authentication for the access method being configured.</p> <p>local: Use the username/password pair configured locally in the switch for the privilege level being configured</p> <p>tacacs: Use a TACACS+ server.</p>
local - or - none	none	n/a	<p>Specifies the secondary (backup) type of authentication being configured.</p> <p>local: The username/password pair configured locally in the switch for the privilege level being configured</p> <p>none: No secondary type of authentication for the specified method/privilege path. (Available only if the primary method of authentication for the access being configured is local.)</p> <p>Note: If you do not specify this parameter in the command line, the switch automatically assigns the secondary method as follows:</p> <ul style="list-style-type: none"> • If the primary method is tacacs, the only secondary method is local. • If the primary method is local, the default secondary method is none.
num-attempts	3	1 - 10	In a given session, specifies how many tries at entering the correct username/password pair are allowed before access is denied and the session terminated.

Configuring the TACACS+ Server for Single Login

In order for the single login feature to work correctly, you need to check some entries in the User Setup on the TACACS+ server.

In the User Setup, scroll to the Advanced TACACS+ Settings section. Make sure the radio button for “Max Privilege for any AAA Client” is checked and the level is set to 15, as shown in Figure 4-4. Privileges are represented by the numbers 0 through 15, with zero allowing only Operator privileges (and requiring two logins) and 15 representing root privileges. The root privilege level is the only level that will allow Manager level access on the switch.



The screenshot shows a configuration window titled "Advanced TACACS+ Settings". Inside, there is a section titled "TACACS+ Enable Control". Below this title are five radio button options, each with a small "x" icon to its right. The first option is "Use Group Level Setting", the second is "No Enable Privilege", the third is "Max Privilege for any AAA Client" (which is selected with a filled radio button), and the fourth is "Define max Privilege on a per network device group basis". Below the third option is a dropdown menu currently showing "Level 15". There is also a small blue icon with a white "x" in the bottom left corner of the settings area.

Figure 4-4. Advanced TACACS+ Settings Section of the TACACS+ Server User Setup

Then scroll down to the section that begins with “Shell” (See Figure 4-5).
Check the Shell box.

Check the Privilege level box and set the privilege level to 15 to allow “root”
privileges. This allows you to use the single login option.

<input checked="" type="checkbox"/>	Shell (exec)	
<input type="checkbox"/>	Access control list	
<input type="checkbox"/>	Auto command	
<input type="checkbox"/>	Callback line	
<input type="checkbox"/>	Callback rotary	
<input type="checkbox"/>	Idle time	
<input type="checkbox"/>	No callback verify	<input type="checkbox"/> Enabled
<input type="checkbox"/>	No escape	<input type="checkbox"/> Enabled
<input type="checkbox"/>	No hangup	<input type="checkbox"/> Enabled
<input checked="" type="checkbox"/>	Privilege level	15
<input type="checkbox"/>	Timeout	123
<input type="checkbox"/>	Custom attributes	

Figure 4-5. The Shell Section of the TACACS+ Server User Setup

As shown in the next table, login and enable access is always available locally through a direct terminal connection to the switch's console port. However, for Telnet access, you can configure TACACS+ to deny access if a TACACS+ server goes down or otherwise becomes unavailable to the switch.

Table 4-2. Primary/Secondary Authentication Table

Access Method and Privilege Level	Authentication Options		Effect on Access Attempts
	Primary	Secondary	
Console — Login	local	none*	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
Console — Enable	local	none	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
Telnet — Login	local	none*	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
	tacacs	none	If Tacacs+ server unavailable, denies access.
Telnet — Enable	local	none	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
	tacacs	none	If Tacacs+ server unavailable, denies access.

Caution Regarding the Use of Local for Login Primary Access

During local authentication (which uses passwords configured in the switch instead of in a TACACS+ server), the switch grants read-only access if you enter the Operator password, and read-write access if you enter the Manager password. For example, if you configure authentication on the switch with Telnet Login Primary as Local and Telnet Enable Primary as Tacacs, when you attempt to Telnet to the switch, you will be prompted for a local password. If you enter the switch's local Manager password (or, if there is no local Manager password configured in the switch) you can bypass the TACACS+ server authentication for Telnet Enable Primary and go directly to read-write (Manager) access. Thus, for either the Telnet or console access method, configuring Login Primary for Local authentication while configuring Enable Primary for TACACS+ authentication is not recommended, as it defeats the purpose of using the TACACS+ authentication. If you want Enable Primary log-in attempts to go to a TACACS+ server, then you should configure both Login Primary and Enable Primary for Tacacs authentication instead of configuring Login Primary to Local authentication.

For example, here is a set of access options and the corresponding commands to configure them:

**Console Login (Operator or Read-Only) Access: Primary using TACACS+ server.
Secondary using Local.**

```
ProCurve (config)# aaa authentication console login tacacs local
```

**Console Enable (Manager or Read/Write) Access: Primary using TACACS+ server.
Secondary using Local.**

```
ProCurve (config)# aaa authentication console enable tacacs local
```

**Telnet Login (Operator or Read-Only) Access: Primary using TACACS+ server.
Secondary using Local.**

```
ProCurve (config)# aaa authentication Telnet login tacacs local
```

**Telnet Enable (Manager or Read/Write Access: Primary using TACACS+ server.
Secondary using Local.**

```
ProCurve (config)# aaa authentication telnet enable tacacs local
```

Deny Access and Close the Session After Failure of Two Consecutive Username/Password Pairs:

```
ProCurve (config)# aaa authentication num-attempts 2
```

Configuring the Switch's TACACS+ Server Access

The `tacacs-server` command configures these parameters:

- **The host IP address(es)** for up to three TACACS+ servers; one first-choice and up to two backups. Designating backup servers provides for a continuation of authentication services in case the switch is unable to contact the first-choice server.
- **An optional encryption key.** This key helps to improve security, and must match the encryption key used in your TACACS+ server application. In some applications, the term “secret key” or “secret” may be used instead of “encryption key”. If you need only one encryption key for the switch to use in all attempts to authenticate through a TACACS+ server, configure a global key. However, if the switch is configured to access multiple TACACS+ servers having different encryption keys, you can configure the switch to use different encryption keys for different TACACS+ servers.
- **The timeout value** in seconds for attempts to contact a TACACS+ server. If the switch sends an authentication request, but does not receive a response within the period specified by the timeout value, the switch resends the request to the next server in its Server IP Addr list, if any. If the switch still fails to receive a response from any TACACS+ server, it reverts to whatever secondary authentication method was configured using the **aaa authentication** command (local or none; see “Configuring the Switch's Authentication Methods” on page 4-11.)

Note

As described under “General Authentication Setup Procedure” on page 4-5, ProCurve recommends that you configure, test, and troubleshoot authentication via Telnet access before you configure authentication via console port access. This helps to prevent accidentally locking yourself out of switch access due to errors or problems in setting up authentication in either the switch or your TACACS+ server.

Syntax: tacacs-server host < ip-addr > [oobm] [key < key-string >]

Adds a TACACS+ server and optionally assigns a server-specific encryption key. The oobm parameter specifies that the operation will go out from the out-of-band management interface. If this parameter is not specified, the operation goes out from the data interface. Refer to Appendix G, “Network Out-of-Band Management” in the Management and Configuration Guide for more information on out-of-band management.

[no] tacacs-server host < ip-addr >

Removes a TACACS+ server assignment (including its server-specific encryption key, if any).

tacacs-server key <key-string>

Enters the optional global encryption key.

[no] tacacs-server key

Removes the optional global encryption key. (Does not affect any server-specific encryption key assignments.)

tacacs-server timeout < 1-255 >

Changes the wait period for a TACACS server response. (Default: 5 seconds.)

Note on Encryption Keys

Encryption keys configured in the switch must exactly match the encryption keys configured in TACACS+ servers the switch will attempt to use for authentication.

If you configure a global encryption key, the switch uses it only with servers for which you have not also configured a server-specific key. Thus, a global key is more useful where the TACACS+ servers you are using all have an identical key, and server-specific keys are necessary where different TACACS+ servers have different keys.

If TACACS+ server “X” does not have an encryption key assigned for the switch, then configuring either a global encryption key or a server-specific key in the switch for server “X” will block authentication support from server “X”.

Name	Default	Range
host <ip-addr> [key <key-string>] [oobm]	none	n/a

TACACS+ Authentication
Configuring TACACS+ on the Switch

Name	Default	Range
------	---------	-------

Specifies the IP address of a device running a TACACS+ server application. Optionally, can also specify the unique, per-server encryption key to use when each assigned server has its own, unique key. For more on the encryption key, see “Using the Encryption Key” on page 4-27 and the documentation provided with your TACACS+ server application. The oobm parameter specifies that the operation will go out from the out-of-band management interface. If this parameter is not specified, the operation goes out from the data interface. Refer to Appendix G, “Network Out-of-Band Management” in the Management and Configuration Guide for more information on out-of-band management.

For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the TACACS+ traffic will go through the out-of-band management (OOBM) port.

You can enter up to three IP addresses; one first-choice and two (optional) backups (one second-choice and one third-choice). Use **show tacacs** to view the current IP address list.
If the first-choice TACACS+ server fails to respond to a request, the switch tries the second address, if any, in the show tacacs list. If the second address also fails, then the switch tries the third address, if any.
(See figure 4-3, “Example of the Switch’s TACACS+ Configuration Listing” on 4-10.)

The priority (first-choice, second-choice, and third-choice) of a TACACS+ server in the switch’s TACACS+ configuration depends on the order in which you enter the server IP addresses:

1. When there are no TACACS+ servers configured, entering a server IP address makes that server the first-choice TACACS+ server.
 2. When there is one TACACS+ server already configured, entering another server IP address makes that server the second-choice (backup) TACACS+ server.
 3. When there are two TACACS+ servers already configured, entering another server IP address makes that server the third-choice (backup) TACACS+ server.
- The above position assignments are fixed. Thus, if you remove one server and replace it with another, the new server assumes the priority position that the removed server had. For example, suppose you configured three servers, A, B, and C, configured in order:
First-Choice: A
Second-Choice: B
Third-Choice: C
 - If you removed server B and then entered server X, the TACACS+ server order of priority would be:
First-Choice: A
Second-Choice: X
Third-Choice: C
 - If there are two or more vacant slots in the TACACS+ server priority list and you enter a new IP address, the new address will take the vacant slot with the highest priority. Thus, if A, B, and C are configured as above and you (1) remove A and B, and (2) enter X and Y (in that order), then the new TACACS+ server priority list would be X, Y, and C.
 - The easiest way to change the order of the TACACS+ servers in the priority list is to remove all server addresses in the list and then re-enter them in order, with the new first-choice server address first, and so on.

To add a new address to the list when there are already three addresses present, you must first remove one of the currently listed addresses.

See also “General Authentication Process Using a TACACS+ Server” on page 4-24.

key <key-string>	none (null) n/a
-------------------------------	-----------------

Name	Default	Range
Specifies the optional, global “encryption key” that is also assigned in the TACACS+ server(s) that the switch will access for authentication. This option is subordinate to any “per-server” encryption keys you assign, and applies only to accessing TACACS+ servers for which you have not given the switch a “per-server” key. (See the host <ip-addr> [key <key-string> entry at the beginning of this table.) For more on the encryption key, see “Using the Encryption Key” on page 4-27 and the documentation provided with your TACACS+ server application.		
timeout <1 - 255>	5 sec	1 - 255 sec
Specifies how long the switch waits for a TACACS+ server to respond to an authentication request. If the switch does not detect a response within the timeout period, it initiates a new request to the next TACACS+ server in the list. If all TACACS+ servers in the list fail to respond within the timeout period, the switch uses either local authentication (if configured) or denies access (if none configured for local authentication).		

Adding, Removing, or Changing the Priority of a TACACS+ Server.

Suppose that the switch was already configured to use TACACS+ servers at 10.28.227.10 and 10.28.227.15. In this case, 10.28.227.15 was entered first, and so is listed as the first-choice server:


```
ProCurve# show tacacs
Status and Counters - TACACS Information
Timeout : 5
Encryption Key : First-Choice TACACS+ Server
```

Server	IP Addr	Closes	Aborts	Errors	Pkts Rx	Pkts Tx
10.28.227.15		0	0	0	0	0
10.28.227.10		0	0	0	0	0

Figure 4-6. Example of the Switch with Two TACACS+ Server Addresses Configured

To move the “first-choice” status from the “15” server to the “10” server, use the **no tacacs-server host <ip-addr>** command to delete both servers, then use **tacacs-server host <ip-addr>** to re-enter the “10” server first, then the “15” server.

The servers would then be listed with the new “first-choice” server, that is:



```
ProCurve# show tacacs
Status and Counters - TACACS Information
Timeout : 5
Encryption Key :
Server IP Addr  Opens    Closes    Aborts    Errors    Pkts Rx  Pkts Tx
-----
10.28.227.10    0         0         0         0         0        0
10.28.227.15    0         0         0         0         0        0
```

Figure 4-7. Example of the Switch After Assigning a Different “First-Choice” Server

To remove the 10.28.227.15 device as a TACACS+ server, you would use this command:

```
ProCurve(config)# no tacacs-server host 10.28.227.15
```

Configuring an Encryption Key. Use an encryption key in the switch if the switch will be requesting authentication from a TACACS+ server that also uses an encryption key. (If the server expects a key, but the switch either does not provide one, or provides an incorrect key, then the authentication attempt will fail.) Use a *global encryption key* if the same key applies to all TACACS+ servers the switch may use for authentication attempts. Use a *per-server encryption key* if different servers the switch may use will have different keys. (For more details on encryption keys, see “Using the Encryption Key” on page 4-27.)

To configure **north01** as a global encryption key:

```
ProCurve(config) tacacs-server key north01
```

To configure **north01** as a per-server encryption key:

```
ProCurve(config)# tacacs-server host 10.28.227.63 key
north01
```

An encryption key can contain up to 100 characters, without spaces, and is likely to be case-sensitive in most TACACS+ server applications.

To delete a global encryption key from the switch, use this command:

```
ProCurve(config)# no tacacs-server key
```

To delete a per-server encryption key in the switch, re-enter the `tacacs-server host` command without the `key` parameter. For example, if you have **north01** configured as the encryption key for a TACACS+ server with an IP address of 10.28.227.104 and you want to eliminate the key, you would use this command:

```
ProCurve(config)# tacacs-server host 10.28.227.104
```

Note

You can save the encryption key in a configuration file by entering this command:

```
Procurve(config)# tacacs-server key <keystring>
```

The *<keystring>* parameter is the encryption key in clear text.

Note

The **show tacacs** command lists the global encryption key, if configured. However, to view any configured per-server encryption keys, you must use **show config** or **show config running** (if you have made TACACS+ configuration changes without executing **write mem**).

Configuring the Timeout Period. The timeout period specifies how long the switch waits for a response to an authentication request from a TACACS+ server before either sending a new request to the next server in the switch's Server IP Address list or using the local authentication option. For example, to change the timeout period from 5 seconds (the default) to 3 seconds:

```
ProCurve(config)# tacacs-server timeout 3
```


How Authentication Operates

General Authentication Process Using a TACACS+ Server

Authentication through a TACACS+ server operates generally as described below. For specific operating details, refer to the documentation you received with your TACACS+ server application.

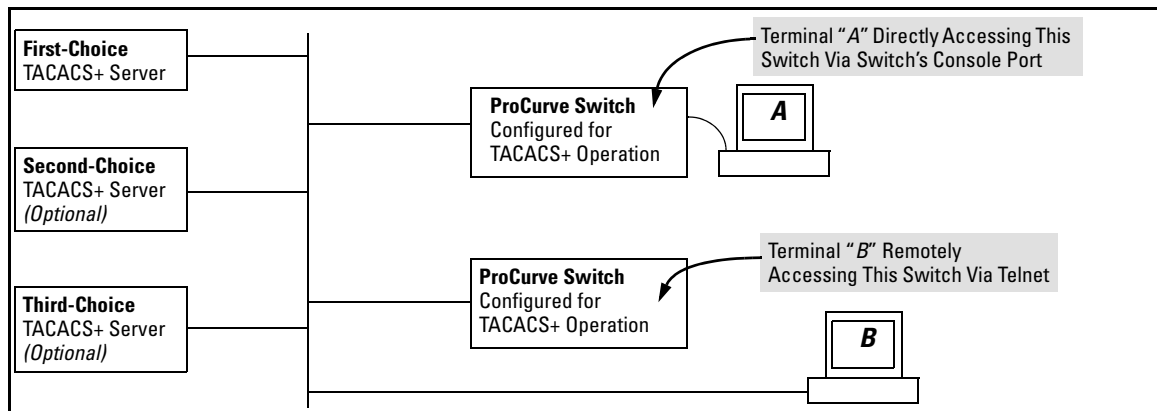


Figure 4-8. Using a TACACS+ Server for Authentication

Using figure 4-8, above, after either switch detects an operator's logon request from a remote or directly connected terminal, the following events occur:

1. The switch queries the first-choice TACACS+ server for authentication of the request.
 - If the switch does not receive a response from the first-choice TACACS+ server, it attempts to query a secondary server. If the switch does not receive a response from any TACACS+ server, then it uses its own local username/password pairs to authenticate the logon request. (See "Local Authentication Process" on page 4-26.)
 - If a TACACS+ server recognizes the switch, it forwards a username prompt to the requesting terminal via the switch.
2. When the requesting terminal responds to the prompt with a username, the switch forwards it to the TACACS+ server.
3. After the server receives the username input, the requesting terminal receives a password prompt from the server via the switch.

4. When the requesting terminal responds to the prompt with a password, the switch forwards it to the TACACS+ server and one of the following actions occurs:
 - If the username/password pair received from the requesting terminal matches a username/password pair previously stored in the server, then the server passes access permission through the switch to the terminal.
 - If the username/password pair entered at the requesting terminal does not match a username/password pair previously stored in the server, access is denied. In this case, the terminal is again prompted to enter a username and repeat steps 2 through 4. In the default configuration, the switch allows up to three attempts to authenticate a login session. If the requesting terminal exhausts the attempt limit without a successful TACACS+ authentication, the login session is terminated and the operator at the requesting terminal must initiate a new session before trying again.

Local Authentication Process

When the switch is configured to use TACACS+, it reverts to local authentication only if one of these two conditions exists:

- “Local” is the authentication option for the access method being used.
- TACACS+ is the primary authentication mode for the access method being used. However, the switch was unable to connect to any TACACS+ servers (or no servers were configured) AND **Local** is the secondary authentication mode being used.

(For a listing of authentication options, see table 4-2, “Primary/Secondary Authentication Table” on 4-16.)

For local authentication, the switch uses the operator-level and manager-level username/password set(s) previously configured locally on the switch. (These are the usernames and passwords you can configure using the CLI password command, the web browser interface, or the menu interface—which enables only local password configuration).

- If the operator at the requesting terminal correctly enters the username/password pair for either access level, access is granted.
- If the username/password pair entered at the requesting terminal does not match either username/password pair previously configured locally in the switch, access is denied. In this case, the terminal is again prompted to enter a username/password pair. In the default configuration, the switch allows up to three attempts. If the requesting terminal exhausts the attempt limit without a successful authentication, the login session is terminated and the operator at the requesting terminal must initiate a new session before trying again.

Note

The switch’s menu allows you to configure only the local Operator and Manager passwords, and not any usernames. In this case, all prompts for local authentication will request only a local password. However, if you use the CLI or the web browser interface to configure usernames for local access, you will see a prompt for both a local username and a local password during local authentication.

Using the Encryption Key

General Operation

When used, the encryption key (sometimes termed “key”, “secret key”, or “secret”) helps to prevent unauthorized intruders on the network from reading username and password information in TACACS+ packets moving between the switch and a TACACS+ server. At the TACACS+ server, a key may include both of the following:

- **Global key:** A general key assignment in the TACACS+ server application that applies to all TACACS-aware devices for which an individual key has not been configured.
- **Server-Specific key:** A unique key assignment in the TACACS+ server application that applies to a specific TACACS-aware device.

Note

Configure a key in the switch only if the TACACS+ server application has this exact same key configured for the switch. That is, if the key parameter in switch “X” does not exactly match the key setting for switch “X” in the TACACS+ server application, then communication between the switch and the TACACS+ server will fail.

Thus, on the TACACS+ server side, you have a choice as to how to implement a key. On the switch side, it is necessary only to enter the key parameter so that it exactly matches its counterpart in the server. For information on how to configure a general or individual key in the TACACS+ server, refer to the documentation you received with the application.

Encryption Options in the Switch

When configured, the encryption key causes the switch to encrypt the TACACS+ packets it sends to the server. When left at “null”, the TACACS+ packets are sent in clear text. The encryption key (or just “key”) you configure in the switch must be identical to the encryption key configured in the corresponding TACACS+ server. If the key is the same for all TACACS+ servers the switch will use for authentication, then configure a global key in the switch. If the key is different for one or more of these servers, use “server-specific” keys in the switch. (If you configure both a global key and one or more per-server keys, the per-server keys will override the global key for the specified servers.)

For example, you would use the next command to configure a global encryption key in the switch to match a key entered as **north40campus** in two target TACACS+ servers. (That is, both servers use the same key for your switch.) Note that you do not need the server IP addresses to configure a global key in the switch:

```
ProCurve(config)# tacacs-server key north40campus
```

Suppose that you subsequently add a third TACACS+ server (with an IP address of 10.28.227.87) that has **south10campus** for an encryption key. Because this key is different than the one used for the two servers in the previous example, you will need to assign a server-specific key in the switch that applies only to the designated server:

```
ProCurve(config)# tacacs-server host 10.28.227.87 key  
south10campus
```

With both of the above keys configured in the switch, the **south10campus** key overrides the **north40campus** key only when the switch tries to access the TACACS+ server having the 10.28.227.87 address.

Controlling Web Browser Interface Access When Using TACACS+ Authentication

Configuring the switch for TACACS+ authentication does not affect web browser interface access. To prevent unauthorized access through the web browser interface, do one or more of the following:

- Configure local authentication (a Manager user name and password and, optionally, an Operator user name and password) on the switch.
- Configure the switch's Authorized IP Manager feature to allow web browser access only from authorized management stations. (The Authorized IP Manager feature does not interfere with TACACS+ operation.)
- Disable web browser access to the switch by going to the System Information screen in the Menu interface and configuring the **Web Agent Enabled** parameter to **No**.

Messages Related to TACACS+ Operation

The switch generates the CLI messages listed below. However, you may see other messages generated in your TACACS+ server application. For information on such messages, refer to the documentation you received with the application.

CLI Message	Meaning
Connecting to Tacacs server	The switch is attempting to contact the TACACS+ server identified in the switch's tacacs-server configuration as the first-choice (or only) TACACS+ server.
Connecting to secondary Tacacs server	The switch was not able to contact the first-choice TACACS+ server, and is now attempting to contact the next (secondary) TACACS+ server identified in the switch's tacacs-server configuration.
Invalid password	The system does not recognize the username or the password or both. Depending on the authentication method (tacacs or local), either the TACACS+ server application did not recognize the username/password pair or the username/password pair did not match the username/password pair configured in the switch.
No Tacacs servers responding	The switch has not been able to contact any designated TACACS+ servers. If this message is followed by the Username prompt, the switch is attempting local authentication.
Not legal combination of authentication methods	For console access , if you select tacacs as the primary authentication method, you must select local as the secondary authentication method. This prevents you from being locked out of the switch if all designated TACACS+ servers are inaccessible to the switch.
Record already exists	When resulting from a tacacs-server host <ip addr> command, indicates an attempt to enter a duplicate TACACS+ server IP address.

Operating Notes

- If you configure Authorized IP Managers on the switch, it is not necessary to include any devices used as TACACS+ servers in the authorized manager list. That is, authentication traffic between a TACACS+ server and the switch is not subject to Authorized IP Manager controls configured on the switch. Also, the switch does not attempt TACACS+ authentication for a management station that the Authorized IP Manager list excludes because, independent of TACACS+, the switch already denies access to such stations.

- When TACACS+ is not enabled on the switch—or when the switch's only designated TACACS+ servers are not accessible—setting a local Operator password without also setting a local Manager password does not protect the switch from manager-level access by unauthorized persons.
- When using the **copy** command to transfer a configuration to a TFTP server, any optional, server-specific and global encryption keys (page 4-18) in the TACACS configuration will not be included in the transferred file. Otherwise, a security breach could occur, allowing access to the TACACS+ username/password information.

RADIUS Authentication, Authorization, and Accounting

Contents

Overview	5-3
Authentication Services	5-3
Accounting Services	5-4
RADIUS-Administered CoS and Rate-Limiting	5-4
RADIUS-Administered Commands Authorization	5-4
SNMP Access to the Switch's Authentication Configuration MIB	5-4
Terminology	5-5
Switch Operating Rules for RADIUS	5-6
General RADIUS Setup Procedure	5-7
Configuring the Switch for RADIUS Authentication	5-8
Outline of the Steps for Configuring RADIUS Authentication	5-9
1. Configure Authentication for the Access Methods You Want RADIUS To Protect	5-10
2. Enable the (Optional) Access Privilege Option	5-13
3. Configure the Switch To Access a RADIUS Server	5-14
4. Configure the Switch's Global RADIUS Parameters	5-17
Using SNMP To View and Configure	
Switch Authentication Features	5-21
Changing and Viewing the SNMP Access Configuration	5-22
Local Authentication Process	5-24
Controlling Web Browser Interface Access	5-25
Commands Authorization	5-26
Enabling Authorization	5-27
Displaying Authorization Information	5-28
Configuring Commands Authorization on a RADIUS Server	5-28
Using Vendor Specific Attributes (VSAs)	5-28

Example Configuration on Cisco Secure ACS for MS Windows	5-30
Example Configuration Using FreeRADIUS	5-32
VLAN Assignment in an Authentication Session	5-34
Tagged and Untagged VLAN Attributes	5-35
Additional RADIUS Attributes	5-36
Configuring RADIUS Accounting	5-37
Operating Rules for RADIUS Accounting	5-39
Steps for Configuring RADIUS Accounting	5-39
1. Configure the Switch To Access a RADIUS Server	5-40
2. Configure Accounting Types and the Controls for Sending Reports to the RADIUS Server	5-42
3. (Optional) Configure Session Blocking and Interim Updating Options	5-44
Viewing RADIUS Statistics	5-46
General RADIUS Statistics	5-46
RADIUS Authentication Statistics	5-48
RADIUS Accounting Statistics	5-49
Changing RADIUS-Server Access Order	5-50
Messages Related to RADIUS Operation	5-53

Overview

Feature	Default	Menu	CLI	Web
Configuring RADIUS Authentication	None	n/a	5-8	n/a
Configuring RADIUS Accounting	None	n/a	5-37	n/a
Configuring RADIUS Authorization	None	n/a	5-26	n/a
Viewing RADIUS Statistics	n/a	n/a	5-46	n/a

RADIUS (*Remote Authentication Dial-In User Service*) enables you to use up to three servers (one primary server and one or two backups) and maintain separate authentication and accounting for each RADIUS server employed. For authentication, this allows a different password for each user instead of having to rely on maintaining and distributing switch-specific passwords to all users. For accounting, this can help you track network resource usage.

Authentication Services

You can use RADIUS to verify user identity for the following types of primary password access to the ProCurve switch:

- Serial port (Console)
- Telnet
- SSH
- SFTP/SCP
- Port-Access (802.1X)

The switch also supports RADIUS accounting for Web Authentication and MAC authentication sessions.

Note

The switch does not support RADIUS security for SNMP (network management) access. For information on blocking access through the web browser interface, refer to “Controlling Web Browser Interface Access” on page 5-25.

Accounting Services

RADIUS accounting on the switch collects resource consumption data and forwards it to the RADIUS server. This data can be used for trend analysis, capacity planning, billing, auditing, and cost analysis.

RADIUS-Administered CoS and Rate-Limiting

The switches covered in this guide take advantage of vendor-specific attributes (VSAs) applied in a RADIUS server to support these optional, RADIUS-assigned attributes:

- 802.1p (CoS) priority assignment to inbound traffic on the specified port(s) (port-access authentication only)
- Per-Port Rate-Limiting on a port with an active link to an authenticated client (port-access authentication only)

RADIUS-Administered Commands Authorization

This feature enables RADIUS server control of an authenticated client’s access to CLI commands on the switch. Refer to “Commands Authorization” on page 5-26.

SNMP Access to the Switch’s Authentication Configuration MIB

The switch’s default configuration allows SNMP access to the hpSwitchAuth MIB (Management Information Base). A management station running an SNMP networked device management application such as ProCurve Manager Plus (PCM+) or HP OpenView can access the switch’s MIB for read access to the switch’s status and read/write access to the switch’s configuration. For more information, including the CLI command to use for disabling this feature, refer to “Using SNMP To View and Configure Switch Authentication Features” on page 5-21.

Terminology

AAA: Authentication, Authorization, and Accounting groups of services provided by the carrying protocol.

CHAP (Challenge-Handshake Authentication Protocol): A challenge-response authentication protocol that uses the Message Digest 5 (MD5) hashing scheme to encrypt a response to a challenge from a RADIUS server.

CoS (Class of Service): Support for priority handling of packets traversing the switch, based on the IEEE 802.1p priority carried by each packet. (For more on this topic, refer to the “Overview” section in the “Quality of Service (QoS)” chapter in the *Advanced Traffic Management Guide* for your switch.)

EAP (Extensible Authentication Protocol): A general PPP authentication protocol that supports multiple authentication mechanisms. A specific authentication mechanism is known as an EAP type, such as MD5-Challenge, Generic Token Card, and TLS (Transport Level Security).

EXEC Session: a service (EXEC shell) granted to the authenticated login user for doing management operations on the ProCurve device.

Host: See **RADIUS Server**.

NAS (Network Access Server): In this case, a ProCurve switch configured for RADIUS security operation.

RADIUS (Remote Authentication Dial In User Service): a protocol for carrying authentication, authorization, and accounting information between a Network Access Server and shared AAA servers in a distributed dial-in networking environment.

RADIUS Client: The device that passes user information to designated RADIUS servers.

RADIUS Host: See RADIUS server.

RADIUS Server: A server running the RADIUS application you are using on your network. This server receives user connection requests from the switch, authenticates users, and then returns all necessary information to the switch. For the ProCurve switch, a RADIUS server can also perform accounting functions. Sometimes termed a *RADIUS host*.

Shared Secret Key: A text value used for encrypting data in RADIUS packets. Both the RADIUS client and the RADIUS server have a copy of the key, and the key is never transmitted across the network.

Vendor-Specific Attribute: A vendor-defined value configured in a RADIUS server to specific an optional switch feature assigned by the server during an authenticated client session.

Switch Operating Rules for RADIUS

- You must have at least one RADIUS server accessible to the switch.
- The switch supports authentication and accounting using up to three RADIUS servers. The switch accesses the servers in the order in which they are listed by **show radius** (page 5-46). If the first server does not respond, the switch tries the next one, and so-on. (To change the order in which the switch accesses RADIUS servers, refer to “Changing RADIUS-Server Access Order” on page 5-50.)
- You can select RADIUS as the primary authentication method for each type of access. (Only one primary and one secondary access method is allowed for each access type.)
- In the ProCurve switch, EAP RADIUS uses MD5 and TLS to encrypt a response to a challenge from a RADIUS server.
- When primary/secondary authentication is set to Radius/Local (for either Login or Enable) and the RADIUS server fails to respond to a client attempt to authenticate, the failure is noted in the Event Log with the message **radius: Can't reach RADIUS server < server-ip-addr >**. When this type of failure occurs, the switch prompts the client again to enter a username and password. In this case, use the local username (if any) and password configured on the switch itself.
- Zero-length usernames or passwords are not allowed for RADIUS authentication, even though allowed by some RADIUS servers.
- TACACS+ is not supported for the web browser interface access.

General RADIUS Setup Procedure

Preparation:

- 1. Configure one to three RADIUS servers to support the switch. (That is, one primary server and one or two backups.) Refer to the documentation provided with the RADIUS server application.
- 2. Before configuring the switch, collect the information outlined below.

Table 5-1. Preparation for Configuring RADIUS on the Switch

- Determine the access methods (console, Telnet, Port-Access (802.1X), web browser interface and/or SSH) for which you want RADIUS as the primary authentication method. Consider both Operator (login) and Manager (enable) levels, as well as which secondary authentication methods to use (local or none) if the RADIUS authentication fails or does not respond.

ProCurve(config)# show authentication

Status and Counters - Authentication Information

Login Attempts : 3

Respect Privilege : Disabled

Access Task	Login Primary	Login Secondary	Enable Primary	Enable Secondary
Console	Radius	Local	Radius	Local
Telnet	Radius	Local	Radius	Local
Port-Access	EapRadius			
Webui	Radius	Local	Radius	Local
SSH	Radius	Local	Radius	Local
Web-Auth	ChapRadius			
MAC-Auth	ChapRadius			

Note: The Webui access task shown in this figure is available only on the switches covered in this guide.

Console access requires Local as secondary method to prevent lockout if the primary RADIUS access fails due to loss of RADIUS server access or other problems with the server.

Figure 5-1. Example of Possible RADIUS Access Assignments

- Determine the IP address(es) of the RADIUS server(s) you want to support the switch. (You can configure the switch for up to three RADIUS servers.)
- If you need to replace the default UDP destination port (1812) the switch uses for authentication requests to a specific RADIUS server, select it before beginning the configuration process.
- If you need to replace the default UDP destination port (1813) the switch uses for accounting requests to a specific Radius server, select it before beginning the configuration process.
- Determine whether you can use one, global encryption key for all RADIUS servers or if unique keys will be required for specific servers. With multiple RADIUS servers, if one key applies to two or more of these servers, then you can configure this key as the global encryption key. For any server whose key differs from the global key you are using, you must configure that key in the same command that you use to designate that server's IP address to the switch.
- Determine an acceptable timeout period for the switch to wait for a server to respond to a request. ProCurve recommends that you begin with the default (five seconds).

- Determine how many times you want the switch to try contacting a RADIUS server before trying another RADIUS server or quitting. (This depends on how many RADIUS servers you have configured the switch to access.)
 - Determine whether you want to bypass a RADIUS server that fails to respond to requests for service. To shorten authentication time, you can set a bypass period in the range of 1 to 1440 minutes for non-responsive servers. This requires that you have multiple RADIUS servers accessible for service requests.
 - Optional: Determine whether the switch access level (Manager or Operator) for authenticated clients can be set by a Service Type value the RADIUS server includes in its authentication message to the switch. (Refer to “2. Enable the (Optional) Access Privilege Option” on page 5-13.)
 - Configure RADIUS on the server(s) used to support authentication on the switch.
-

Configuring the Switch for RADIUS Authentication

RADIUS Authentication Commands	Page
aaa authentication	5-10
console telnet ssh web < enable login <local radius>> web-based mac-based <chap-radius peap-radius>	5-10
[local none authorized]	5-10
[login privilege-mode]*	5-13
[no] radius-server host < IP-address >	5-14
[auth-port < port-number >]	5-14
[acct-port < port-number >]	5-14, 5-40
[key < server-specific key-string >]	5-14
[no] radius-server key < global key-string >	5-18
radius-server timeout < 1 - 15>	5-18
radius-server retransmit < 1 - 5 >	5-18
[no] radius-server dead-time < 1 - 1440 >	5-19
show radius	5-46
[< host < ip-address>]	5-47
show authentication	5-48
show radius authentication	5-49

*The **web** authentication option for the web browser interface is available on the switches covered in this guide.

Outline of the Steps for Configuring RADIUS Authentication

There are three main steps to configuring RADIUS authentication:

1. Configure RADIUS authentication for controlling access through one or more of the following
 - Serial port
 - Telnet
 - SSH
 - Port-Access (802.1X)
 - Web browser interface
2. Enable RADIUS authentication on the switch to override the default authentication operation of automatically assigning an authenticated client to the Operator privilege level. This optional feature applies the privilege level specified by the Service Type value received from the RADIUS server. (Refer to “1. Configure Authentication for the Access Methods You Want RADIUS To Protect” on page 5-10.)
3. Configure the switch for accessing one or more RADIUS servers (one primary server and up to two backup servers):

Note

This step assumes you have already configured the RADIUS server(s) to support the switch. Refer to the documentation provided with the RADIUS server documentation.)

- Server IP address
 - (Optional) UDP destination port for authentication requests (default: 1812; recommended)
 - (Optional) UDP destination port for accounting requests (default: 1813; recommended)
 - (Optional) encryption key for use during authentication sessions with a RADIUS server. This key overrides the global encryption key you can also configure on the switch, and must match the encryption key used on the specified RADIUS server. (Default: null)
4. Configure the global RADIUS parameters.
 - **Server Key:** This key must match the encryption key used on the RADIUS servers the switch contacts for authentication and accounting services unless you configure one or more per-server keys. (Default: null.)

- **Timeout Period:** The timeout period the switch waits for a RADIUS server to reply. (Default: 5 seconds; range: 1 to 15 seconds.)
- **Retransmit Attempts:** The number of retries when there is no server response to a RADIUS authentication request. (Default: 3; range of 1 to 5.)
- **Server Dead-Time:** The period during which the switch will not send new authentication requests to a RADIUS server that has failed to respond to a previous request. This avoids a wait for a request to time out on a server that is unavailable. If you want to use this feature, select a dead-time period of 1 to 1440 minutes. (Default: 0—disabled; range: 1 - 1440 minutes.) If your first-choice server was initially unavailable, but then becomes available before the dead-time expires, you can nullify the dead-time by resetting it to zero and then trying to log on again. As an alternative, you can reboot the switch, (thus resetting the dead-time counter to assume the server is available) and then try to log on again.
- **Number of Login Attempts:** This is actually an **aaa authentication** command. It controls how many times per session a RADIUS client (and clients using other forms of access) can try to log in with the correct username and password. (Default: Three times per session.)

(For RADIUS accounting features, refer to “Configuring RADIUS Accounting” on page 5-37.)

1. Configure Authentication for the Access Methods You Want RADIUS To Protect

This section describes how to configure the switch for RADIUS authentication through the following access methods:

- **Console:** Either direct serial-port connection or modem connection.
- **Telnet:** Inbound Telnet must be enabled (the default).
- **SSH:** To use RADIUS for SSH access, first configure the switch for SSH operation. Refer to chapter 6, “Configuring Secure Shell (SSH)”.
- **Web:** You can enable RADIUS authentication for web browser interface access to the switch.

You can configure RADIUS as the primary password authentication method for the above access methods. You also need to select either **local**, **none**, or **authorized** as a secondary, or backup, method. Note that for console access, if you configure **radius** (or **tacacs**) for primary authentication, you must config-

ure **local** for the secondary method. This prevents the possibility of being completely locked out of the switch in the event that all primary access methods fail.

Syntax: `aaa authentication < console | telnet | ssh | web | < enable | login <local | radius>> web-based | mac-based <chap-radius | peap-radius>>`

Configures RADIUS as the primary password authentication method for console, Telnet, SSH, and/or the web browser interface. (The default primary < enable | login > authentication is local.)

`<console | telnet | ssh | web>`

`[< local | none | authorized >]`

*Provides options for secondary authentication (default: **none**). Note that for console access, secondary authentication must be **local** if primary access is not **local**. This prevents you from being locked out of the switch in the event of a failure in other access methods.*

`<<web-based | mac-based> login> <chap-radius | peap-mschapv2>:`
*Password authentication for web-based or mac-based port access to the switch. Use **peap-mschapv2** when you want password verification without requiring access to a plain text password; it is more secure.*

Default: **chap-radius**

`[none | authorized];` *Provides options for secondary authentication. The **none** option specifies that a backup authentication method is not used. The **authorized** option allows access without authentication. (default: **none**).*

In certain situations, RADIUS servers can become isolated from the network. Users are not able to access the network resources configured with RADIUS access protection and are rejected. To address this situation, configuring the **authorized** secondary authentication method allows users unconditional access to the network when the primary authentication method fails because the RADIUS servers are unreachable.

Caution

Configuring **authorized** as the secondary authentication method used when there is a failure accessing the RADIUS servers allows clients to access the network unconditionally. Use this method with care.

Figure 5-2 shows an example of the **show authentication** command displaying **authorized** as the secondary authentication method for port-access, Web-auth access, and MAC-auth access. Since the configuration of **authorized** means no authentication will be performed and the client has unconditional access to the network, the “Enable Primary” and “Enable Secondary” fields are not applicable (N/A).

```
ProCurve(config)# show authentication
```

Status and Counters - Authentication Information

Login Attempts : 3
Respect Privilege : Disabled

Access Task	Login Primary	Login Secondary	Enable Primary	Enable Secondary
Console	Local	None	Local	None
Telnet	Local	None	Local	None
Port-Access	Local	Authorized	N/A	N/A
Webui	Local	None	Local	None
SSH	Local	None	Local	None
Web-Auth	ChapRadius	Authorized	N/A	N/A
MAC-Auth	ChapRadius	Authorized	N/A	N/A

The access methods with secondary authentication configured as **authorized** allows the client access to the network even if the RADIUS server is unreachable.

Figure 5-2. Example of AAA Authentication Using Authorized for the Secondary Authentication Method

Suppose you already configured local passwords on the switch, but want RADIUS to protect primary Telnet and SSH access without allowing a secondary Telnet or SSH access option (the switch’s local passwords):

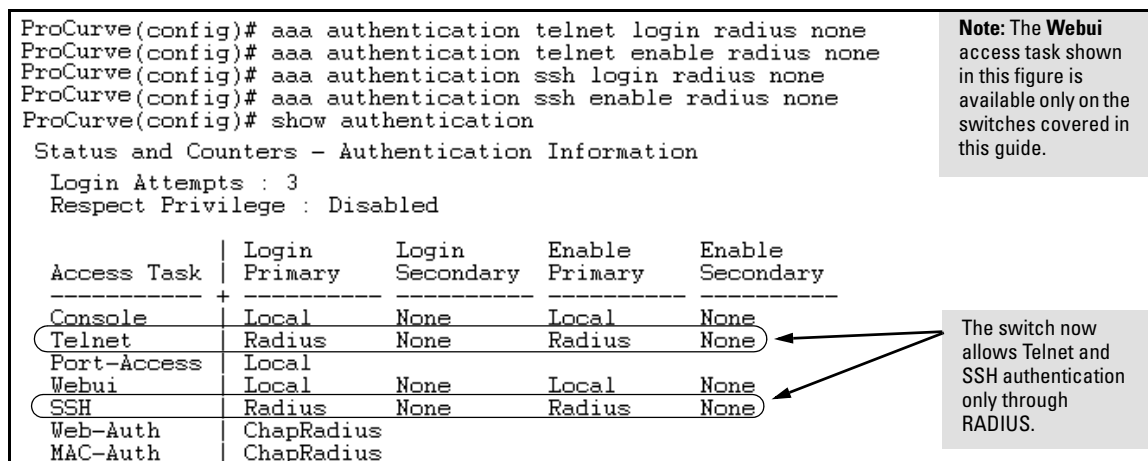


Figure 5-3. Example Configuration for RADIUS Authentication

Note

If you configure the Login Primary method as **local** instead of **radius** (and local passwords are configured on the switch), then clients connected to your network can gain access to either the Operator or Manager level without encountering the RADIUS authentication specified for Enable Primary. Refer to “Local Authentication Process” on page 5-24.

2. Enable the (Optional) Access Privilege Option

In the default RADIUS operation, the switch automatically admits any authenticated client to the Login (Operator) privilege level, even if the RADIUS server specifies Enable (Manager) access for that client. Thus, an authenticated user authorized for the Manager privilege level must authenticate again to change privilege levels. Using the optional **login privilege-mode** command overrides

this default behavior for clients with Enable (manager) access. That is, with **privilege-mode** enabled, the switch immediately allows Enable (Manager) access to a client for whom the RADIUS server specifies this access level.

Syntax: [no] aaa authentication login privilege-mode

When enabled, the switch reads the Service-Type field in the client authentication received from a RADIUS server. The following table describes the applicable Service-Type values and corresponding client access levels the switch allows upon authentication by the server.

Service-Type	Value	Client Access Level
Administrative-User	6	Manager
NAS-Prompt-User	7	Operator
Any Other Type	Any Value Except 6 or 7	Access Denied

This feature applies to console (serial port), Telnet, SSH, and web browser interface access to the switch. It does not apply to 802.1X port-access.

Notes: *While this option is enabled, a Service-Type value other than 6 or 7, or an unconfigured (null) Service-Type causes the switch to deny access to the requesting client.*

*The **no** form of the command returns the switch to the default RADIUS authentication operation. The default behavior for most interfaces is that a client authorized by the RADIUS server for Enable (Manager) access will be prompted twice, once for Login (Operator) access and once for Enable access. In the default RADIUS authentication operation, the switch's web browser interface requires only one successful authentication request. For more information on configuring the Service Type in your RADIUS application, refer to the documentation provided with the application.*

3. Configure the Switch To Access a RADIUS Server

This section describes how to configure the switch to interact with a RADIUS server for both authentication and accounting services.

Note

If you want to configure RADIUS accounting on the switch, go to page 5-37: “Configuring RADIUS Accounting” instead of continuing here.

Syntax: [no] radius-server host < ip-address > [oobm]

*Adds a server to the RADIUS configuration or (with **no**) deletes a server from the configuration. You can configure up to three RADIUS server addresses. The switch uses the first server it successfully accesses. (Refer to “Changing the RADIUS Server Access Order” on page 5-50.)*

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the RADIUS traffic will go through the out-of-band management (OOBM) port.*

[auth-port < port-number >]

*Optional. Changes the UDP destination port for authentication requests to the specified RADIUS server (host). If you do not use this option with the **radius-server host** command, the switch automatically assigns the default authentication port number. The **auth-port** number must match its server counterpart. (Default: **1812**)*

[acct-port < port-number >]

*Optional. Changes the UDP destination port for accounting requests to the specified RADIUS server. If you do not use this option with the **radius-server host** command, the switch automatically assigns the default accounting port number. The **acct-port** number must match its server counterpart. (Default: **1813**)*

[key < key-string >]

Optional. Specifies an encryption key for use during authentication (or accounting) sessions with the specified server. This key must match the encryption key used on the RADIUS server. Use this command only if the specified server requires a different encryption key than configured for the global encryption key.

Note: *Formerly, when you saved the configuration file using Xmodem (Xmodem is supported over OA, Onboard Administrator, but not over USB) or TFTP, the RADIUS encryption key information was not saved in the file. This caused RADIUS authentication to break when the startup configuration file was loaded back onto the switch. You now can save the configured RADIUS shared secret (encryption) key to a configuration file by entering the following commands:*

include-credentials
write memory

For more information, see “Saving Security Credentials in a Config File” on page 2-10 in this guide.

no radius-server host < ip-address > key

*Use the **no** form of the command to remove the key for a specified server.*

For example, suppose you have configured the switch as shown in figure 5-4 and you now need to make the following changes:

1. Change the encryption key for the server at 10.33.18.127 to “source0127”.
2. Add a RADIUS server with an IP address of 10.33.18.119 and a server-specific encryption key of “source0119”.

```
ProCurve# show radius
Status and Counters - General RADIUS Information
  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 5
  Global Encryption Key :

  Server IP Addr  Auth  Acct
  -----
  10.33.18.127    1812  1813  TempKey01
```

Figure 5-4. Sample Configuration for RADIUS Server Before Changing the Key and Adding Another Server

To make the changes listed prior to figure 5-4, you would do the following:

```
ProCurve(config)# radius-server host 10.33.18.127 key source0127
ProCurve(config)# radius-server host 10.33.18.119 key source0119
ProCurve(config)# show radius
Status and Counters - General RADIUS Information
  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 5
  Global Encryption Key :

  Server IP Addr  Auth  Acct
  -----
  10.33.18.127    1812  1813  source0127
  10.33.18.119    1812  1813  source0119
```

Changes the key for the existing server to "source0127" (step 1, above).

Adds the new RADIUS server with its required "source0119" key.

Lists the switch's new RADIUS server configuration. Compare this with

Figure 5-5. Sample Configuration for RADIUS Server After Changing the Key and Adding Another Server

To change the order in which the switch accesses RADIUS servers, refer to "Changing RADIUS-Server Access Order" on page 5-50.

4. Configure the Switch's Global RADIUS Parameters

You can configure the switch for the following global RADIUS parameters:

- **Number of login attempts:** In a given session, specifies how many tries at entering the correct username and password pair are allowed before access is denied and the session terminated. (This is a general **aaa authentication** parameter and is not specific to RADIUS.)

- **Global server key:** The server key the switch will use for contacts with all RADIUS servers for which there is not a server-specific key configured by **radius-server host < ip-address > key < key-string >**. This key is optional if you configure a server-specific key for each RADIUS server entered in the switch. (Refer to “3. Configure the Switch To Access a RADIUS Server” on page 5-14.)
- **Server timeout:** Defines the time period in seconds for authentication attempts. If the timeout period expires before a response is received, the attempt fails.
- **Server dead time:** Specifies the time in minutes during which the switch avoids requesting authentication from a server that has not responded to previous requests.
- **Retransmit attempts:** If the first attempt to contact a RADIUS server fails, specifies how many retries you want the switch to attempt on that server.

Syntax: aaa authentication num-attempts < 1 - 10 >

Specifies how many tries for entering the correct user-name and password before shutting down the session due to input errors. (Default: 3; Range: 1 - 10).

[no] radius-server

key < global-key-string >

Specifies the global encryption key the switch uses with servers for which the switch does not have a server-specific key assignment. This key is optional if all RADIUS server addresses configured in the switch include a server-specific encryption key. (Default: Null.)

dead-time < 1 - 1440 >

Optional. Specifies the time in minutes during which the switch will not attempt to use a RADIUS server that has not responded to an earlier authentication attempt. (Default: 0; Range: 1 - 1440 minutes)

radius-server timeout < 1 - 15 >

Specifies the maximum time the switch waits for a response to an authentication request before counting the attempt as a failure. (Default: 3 seconds; Range: 1 - 15 seconds)

radius-server retransmit < 1 - 5 >

If a RADIUS server fails to respond to an authentication request, specifies how many retries to attempt before closing the session. Default: 3; Range: 1 - 5)

Note

Where the switch has multiple RADIUS servers configured to support authentication requests, if the first server fails to respond, then the switch tries the next server in the list, and so on. If none of the servers respond, then the switch attempts to use the secondary authentication method configured for the type of access being attempted (console, Telnet, or SSH). If this occurs, refer to “RADIUS-Related Problems” in the Troubleshooting chapter of the Management and Configuration Guide for your switch.

For example, suppose that your switch is configured to use three RADIUS servers for authenticating access through Telnet and SSH. Two of these servers use the same encryption key. In this case your plan is to configure the switch with the following global authentication parameters:

- Allow only two tries to correctly enter username and password.
- Use the global encryption key to support the two servers that use the same key. (For this example, assume that you did not configure these two servers with a server-specific key.)
- Use a dead-time of five minutes for a server that fails to respond to an authentication request.
- Allow three seconds for request timeouts.
- Allow two retries following a request that did not receive a response.

```
ProCurve(config)# aaa authentication num-attempts 2
ProCurve(config)# radius-server key My-Global-Key-1099
ProCurve(config)# radius-server dead-time 5
ProCurve(config)# radius-server timeout 3
ProCurve(config)# radius-server retransmit 2
ProCurve(config)# write mem
```

Figure 5-6. Example of Global Configuration Exercise for RADIUS Authentication

RADIUS Authentication, Authorization, and Accounting

Configuring the Switch for RADIUS Authentication

```
ProCurve(config)# show authentication
```

Status and Counters - Authentication Information

Login Attempts : 2
Respect Privilege : Disabled

Access Task	Login Primary	Login Secondary	Enable Primary	Enable Secondary
Console	Local	None	Local	None
Telnet	Radius	None	Radius	None
Port-Access	Local			
Webui	Local	None	Local	None
SSH	Radius	None	Radius	None
Web-Auth	ChapRadius			
MAC-Auth	ChapRadius			

ProCurve(config)# show radius

Status and Counters - General RADIUS Information

Deadtime(min) : 5
Timeout(secs) : 3
Retransmit Attempts : 2
Global Encryption Key : My-Global-Key-1099

Server IP Addr	Auth Port	Acct Port	Encryption Key
10.33.18.127	1812	1813	source0127
10.33.18.119	1812	1813	
10.33.18.151	1812	1813	

Note: The **Webui** access task shown in this figure is available only on the switches covered in this guide.

After two attempts failing due to username or password entry errors, the switch will terminate the session.

Global RADIUS parameters from figure 5-6.

Server-specific encryption key for the RADIUS server that will not use the global encryption key.

These two servers will use the global encryption key.

Figure 5-7. Listings of Global RADIUS Parameters Configured In Figure 5-6

Using SNMP To View and Configure Switch Authentication Features

SNMP MIB object access is available for switch authentication configuration (hpSwitchAuth) features. This means that the switches covered by this Guide allow, by default, manager-only SNMP read/write access to a subset of the authentication MIB objects for the following features:

- number of primary and secondary login and enable attempts
- TACACS+ server configuration and status
- RADIUS server configuration
- selected 802.1X settings
- key management subsystem chain configuration
- key management subsystem key configuration
- local switch operator and manager usernames and passwords

With SNMP access to the hpSwitchAuth MIB enabled, a device with management access to the switch can view the configuration for the authentication features listed above (excluding usernames, passwords, and keys). Using SNMP sets, a management device can change the authentication configuration (*including* changes to usernames, passwords, and keys). Operator read/write access to the authentication MIB is always denied.

Security Notes

All usernames, passwords, and keys configured in the hpSwitchAuth MIB are not returned via SNMP, and the response to SNMP queries for such information is a null string. However, SNMP sets can be used to configure username, password, and key MIB objects.

To help prevent unauthorized access to the switch's authentication MIB, ProCurve recommends following the "SNMP Security Guidelines" on page 1-15.

If you do not want to use SNMP access to the switch's authentication configuration MIB, then use the **snmp-server mib hpswitchauthmib excluded** command to disable this access, as described in the next section.

If you choose to leave SNMP access to the security MIB open (the default setting), ProCurve recommends that you configure the switch with the SNMP version 3 management and access security feature, and disable SNMP version 2c access. (Refer to "Access Security Features" on page 1-3.)

Changing and Viewing the SNMP Access Configuration

Syntax: `snmp-server mib hpswitchauthmib < excluded | included >`

included: *Enables manager-level SNMP read/write access to the switch's authentication configuration (hpSwitchAuth) MIB.*

excluded: *Disables manager-level SNMP read/write access to the switch's authentication configuration (hpSwitchAuth) MIB.*
(Default: **included**)

Syntax: `show snmp-server`

*The output for this command has been enhanced to display the current access status of the switch's authentication configuration MIB in the **Excluded MIBs** field.*

For example, to disable SNMP access to the switch's authentication MIB and then display the result in the Excluded MIB field, you would execute the following two commands.

```
ProCurve(config)# snmp-server mib hpswitchauthmib excluded
ProCurve(config)# show snmp-server
```

SNMP Communities

Community Name	MIB View	Write Access
public	Manager	Unrestricted

Trap Receivers

Link-Change Traps Enabled on Ports [All] : All

Send Authentication Traps [No] : No

Address	Community	Events Sent in Trap
---------	-----------	---------------------

Excluded MIBs

hpSwitchAuthenticationMIB

This command disables SNMP security MIB access.

Indicates that SNMP security MIB access is disabled, which is the nondefault setting.

Figure 5-8. Disabling SNMP Access to the Authentication MIB and Displaying the Result

An alternate method of determining the current Authentication MIB access state is to use the **show run** command.

```
ProCurve(config)# show run

Running configuration:

; 498358-B21 Configuration Editor; Created on release #Z.14.XX

hostname "ProCurve"
[snmp-server mib hpSwitchAuthMIB excluded] ← Indicates that SNMP access to the
ip default-gateway 10.10.24.55             authentication configuration MIB
snmp-server community "public" Operator    (hpSwitchAuth) is disabled.
vlan 1
    name "DEFAULT_VLAN"
    untagged A1-A24,B1-B4
    ip address 10.10.24.100 255.255.255.0
    exit
password manager
```

Figure 5-9. Using the show run Command to View the Current Authentication MIB Access State

Local Authentication Process

When the switch is configured to use RADIUS, it reverts to local authentication only if one of these two conditions exists:

- **Local** is the authentication option for the access method being used.
- The switch has been configured to query one or more RADIUS servers for a primary authentication request, but has not received a response, and **Local** is the configured secondary option.

For local authentication, the switch uses the Operator-level and Manager-level username/password set(s) previously configured locally on the switch. (These are the usernames and passwords you can configure using the CLI password command, the web browser interface, or the menu interface—which enables only local password configuration).

- If the operator at the requesting terminal correctly enters the username/password pair for either access level (Operator or Manager), access is granted on the basis of which username/password pair was used. For example, suppose you configure Telnet primary access for RADIUS and Telnet secondary access for local. If a RADIUS access attempt fails, then you can still get access to either the Operator or Manager level of the switch by entering the correct username/password pair for the level you want to enter.
- If the username/password pair entered at the requesting terminal does not match either local username/password pair previously configured in the switch, access is denied. In this case, the terminal is again prompted to enter a username/password pair. In the default configuration, the switch allows up to three attempts. If the requesting terminal exhausts the attempt limit without a successful authentication, the login session is terminated and the operator at the requesting terminal must initiate a new session before trying again.

Controlling Web Browser Interface Access

To help prevent unauthorized access through the web browser interface, do one or more of the following:

- Configure the switch to support RADIUS authentication for web browser interface access (Web Authentication, Chapter 7).
- Options for the switches covered in this guide:
 - Configure local authentication (a Manager user name and password and, optionally, an Operator user name and password) on the switch.
 - Configure the switch's Authorized IP Manager feature to allow web browser access only from authorized management stations. (The Authorized IP Manager feature does not interfere with TACACS+ operation.)
 - Use one of the following methods to disable web browser access to the switch via http (Port 80):

CLI: **no web-management**

Menu Interface—From the Main menu, select the following:

2. Switch Configuration

1. System Information

Web Agent Enabled: No

Commands Authorization

The RADIUS protocol combines user authentication and authorization steps into one phase. The user must be successfully authenticated before the RADIUS server will send authorization information (from the user's profile) to the Network Access Server (NAS). After user authentication has occurred, the authorization information provided by the RADIUS server is stored on the NAS for the duration of the user's session. Changes in the user's authorization profile during this time will not be effective until after the next authentication occurs.

You can limit the services for a user by enabling AAA RADIUS authorization. The NAS uses the information set up on the RADIUS server to control the user's access to CLI commands.

The authorization type implemented on the switches covered in this guide is the "commands" method. This method explicitly specifies on the RADIUS server which commands are allowed on the client device for authenticated users. This is done on a per-user or per-group basis.

Note

The commands authorization will only be executed for commands entered from Telnet, SSH, or console sessions. The Web management interface is not supported.

By default, all users may execute a minimal set of commands regardless of their authorization status, for example, "exit" and "logout". This minimal set of commands can prevent deadlock on the switch due to an error in the user's authorization profile on the RADIUS server.

Enabling Authorization

To configure authorization for controlling access to the CLI commands, enter this command at the CLI.

Syntax: [no] aaa authorization <commands> <radius | none>

Configures authorization for controlling access to CLI commands. When enabled, the switch checks the list of commands supplied by the RADIUS server during user authentication to determine if a command entered by the user can be executed.

radius: *The NAS requests authorization information from the RADIUS server. Authorization rights are assigned by user or group.*

none: *The NAS does not request authorization information.*

For example, to enable the RADIUS protocol as the authorization method:

```
ProCurve(config)# aaa authorization commands radius
```

When the NAS sends the RADIUS server a valid username and password, the RADIUS server sends an Access-Accept packet that contains two attributes—the command list and the command exception flag. When an authenticated user enters a command on the switch, the switch examines the list of commands delivered in the RADIUS Access-Accept packet as well as the command exception flag, which indicates whether the user has permission to execute the commands in the list. See *Configuring the RADIUS Server* on page 5-28.

After the Access-Accept packet is delivered, the command list resides on the switch. Any changes to the user's command list on the RADIUS server are not seen until the user is authenticated again.

Displaying Authorization Information

You can show the authorization information by entering this command:

Syntax: show authorization

Configures authorization for controlling access to CLI commands. When enabled, the switch checks the list of commands supplied by the RADIUS server during user authentication to determine if a command entered by the user can be executed.

An example of the output is shown.

```
ProCurve(config)# show authorization

Status and Counters - Authorization Information

Type      | Method
-----+-----
Commands  | RADIUS
```

Figure 5-10. Example of Show Authorization Command

Configuring Commands Authorization on a RADIUS Server

Using Vendor Specific Attributes (VSAs)

Some RADIUS-based features implemented on ProCurve switches use HP VSAs for information exchange with the RADIUS server. RADIUS Access-Accept packets sent to the switch may contain the vendor-specific information. The attributes supported with **commands** authorization are:

- **HP-Command-String:** List of commands (regular expressions) that are permitted (or denied) execution by the user. The commands are delimited by semi-colons and must be between 1 and 249 characters in length. Multiple instances of this attribute may be present in Access-Accept packets. (A single instance may be present in Accounting-Request packets.)
- **HP-Command-Exception:** A flag that specifies whether the commands indicated by the HP-Command-String attribute are permitted or denied to the user. A zero (0) means permit all listed commands and deny all others; a one (1) means deny all listed commands and permit all others.

The results of using the HP-Command-String and HP-Command-Exception attributes in various combinations are shown below.

HP-Command-String	HP-Command-Exception	Description
Not present	Not present	If command authorization is enabled and the RADIUS server does not provide any authorization attributes in an Access-Accept packet, the user is denied access to the server. This message appears: "Access denied: no user's authorization info supplied by the RADIUS server."
Not present	DenyList-PermitOthers(1)	Authenticated user is allowed to execute all commands available on the switch.
Not present	PermitList-DenyOthers(0)	Authenticated user can only execute a minimal set of commands (those that are available by default to any user).
Commands List	DenyList-PermitOthers(1)	Authenticated user may execute all commands except those in the Commands list.
Commands List	PermitList-DenyOthers(0)	Authenticated user can execute only those commands provided in the Commands List, plus the default commands.
Commands List	Not present	Authenticated user can only execute commands from the Commands List, plus the default commands.
Empty Commands List	Not present	Authenticate user can only execute a minimal set of commands (those that are available by default to any user).
Empty Commands List	DenyList-PermitOthers(1)	Authenticated user is allowed to execute all commands available on the switch.
Empty Commands List	PermitList-DenyOthers(0)	Authenticate user can only execute a minimal set of commands (those that are available by default to any user).

You must configure the RADIUS server to provide support for the HP VSAs. There are multiple RADIUS server applications; the two examples below show how a dictionary file can be created to define the VSAs for that RADIUS server application.

Example Configuration on Cisco Secure ACS for MS Windows

It is necessary to create a dictionary file that defines the VSAs so that the RADIUS server application can determine which VSAs to add to its user interface. The VSAs will appear below the standard attributes that can be configured in the application.

The dictionary file must be placed in the proper directory on the RADIUS server. Follow these steps.

1. Create a dictionary file (for example, hp.ini) containing the HP VSA definitions, as shown in the example below.

```
; [User Defined Vendor]
;
; The Name and IETF vendor code and any VSAs MUST be unique.
;
; One or more VSAs named (max 255)
;
; Each named VSA requires a definition section...
;
; Types are STRING, INTEGER, IPADDR
;
; The profile specifies usage, IN for accounting, OUT for
  authorization,
; MULTI if more than a single instance is allowed per
  RADIUS message.
; Combinations are allowed, e.g. "IN", "MULTI OUT",
  "MULT IN OUT"
;
; Enumerations are optional for INTEGER attribute types

[User Defined Vendor]

Name=HP
IETF Code=11
VSA 2=Hp-Command-String
VSA 3=Hp-Command-Exception

[Hp-Command-String]

Type=STRING
Profile=IN OUT

[Hp-Command-Exception]

Type=INTEGER
```

Profile=IN OUT

Enums=Hp-Command-Exception-Types

[Hp-Command-Exception-Types]

0=PermitList

1=DenyList

2. Copy the hp.ini dictionary file to c:\program files\cisco acs 3.2\utils (or the \utils directory wherever acs is installed).
3. From the command prompt execute the following command:

```
c:\Program files\CiscoSecure ACS v3.2\utils>  
csutil -addudv 0 hp.ini
```

The zero (0) is the slot number. You will see some processing messages:

Adding or removing vendors requires ACS services to be re-started. Please make sure regedit is not running as it can prevent registry backup/restore operations.

Are you sure you want to proceed? (Y or N) y

Parsing [.\hp.ini] for addition at UDV slot [0]

Stopping any running services

Creating backup of current config

Adding Vendor [HP} added as [RADIUS (HP)]

Done

Checking new configuration...

New configuration OK

Re-starting stopped services

4. Start the registry editor (regedit) and browse to HKEY_LOCAL_MACHINE\software\cisco\CiscoAAA v3.2\NAS Vendors tree.

Cisco adds the entry into this tree for each custom vendor. The id is 100 + the slot number used in the previous command (100 + 0, as it was added in slot 0). Look in the key to verify the vendor name and id.

5. Go to:

HKEY_LOCAL_MACHINE\software\cisco\CiscoAAAv3.2\
CSRradius\ExtensionPoints\002\AssociatedWithVendors

6. Right click and then select **New > key**. Add the vendor Id number that you determined in step 4 (100 in the example).
7. Restart all Cisco services.
8. The newly created HP RADIUS VSA appears only when you configure an AAA client (NAS) to use the HP VSA RADIUS attributes. Select Network Configuration and add (or modify) an AAA entry. In the Authenticate Using field choose RADIUS(HP) as an option for the type of security control protocol.
9. Select **Submit + Restart** to effect the change. The HP RADIUS VSA attributes will appear in Cisco ACS configurations, for example, "Interface Configuration", "Group Setup", "User Setup".

To enable the processing of the HP-Command-String VSA for RADIUS accounting:

1. Select **System Configuration**.
2. Select **Logging**.
3. Select **CSV RADIUS Accounting**. In the Select Columns to Log section, add the HP-Command-String attribute to the Logged Attributes list.
4. Select **Submit**.
5. Select **Network Configuration**. In the AAA Clients section, select an entry in the AAA Client Hostname column. You will go to the AAA Client Setup screen.
6. Check the box for **Log Update/Watchdog Packets from this AAA Client**.
7. Click **Submit + Restart**. You should be able to see the HP-Command-String attribute in the RADIUS accounting reports.

You can enter the commands you wish to allow or deny with the special characters used in standard regular expressions (c, ., \, [list], [^list], *, ^, \$). Commands must be between 1-249 characters in length.

Example Configuration Using FreeRADIUS

1. Create a dictionary file (for example, dictionary.hp) containing HP VSA definitions. An example file is:

```
#
# dictionary.hp
#
# As posted to the list by User <user_email>
#
# Version: $Id: dictionary.hp, v 1.0 2006/02/23 17:07:07
#

VENDOR            Hp            11

# HP Extensions

ATTRIBUTE          Hp-Command-String      2      string      Hp
ATTRIBUTE          Hp-Command-Exception   3      integer     Hp

# Hp-Command-Exception Attribute Values

VALUE              Hp-Command-Exception    Permit-List      0
VALUE              Hp-Command-Exception    Deny-List        1
```

2. Find the location of the dictionary files used by FreeRADIUS (try /usr/local/share/freeradius).
3. Copy dictionary.hp to that location. Open the existing dictionary file and add this entry:
\$ INCLUDE dictionary.hp
4. You can now use HP VSAs with other attributes when configuring user entries.

VLAN Assignment in an Authentication Session

A switch supports concurrent 802.1X and either Web- or MAC-authentication sessions on a port (with up to 32 clients allowed). If you have configured RADIUS as the primary authentication method for a type of access, when a client authenticates on a port, the RADIUS server assigns an untagged VLAN that is statically configured on the switch for use in the authentication session. (For information on how to configure a user profile on a RADIUS server with the VLAN to be assigned for 802.1X, Web, or MAC authentication, refer to the documentation provided with the RADIUS server application.)

If a switch port is configured to accept multiple 802.1X and/or Web- or MAC-Authentication client sessions, all authenticated clients must use the same port-based, untagged VLAN membership assigned for the earliest, currently active client session. On a port where one or more authenticated client sessions are already running, all clients are on the same untagged VLAN. If the RADIUS server subsequently authenticates a new client, but attempts to re-assign the port to a different, untagged VLAN than the one already in use for the previously existing, authenticated client sessions, the connection for the new client will fail.

Tagged and Untagged VLAN Attributes

When you configure a user profile on a RADIUS server to assign a VLAN to an authenticated client, you can use either the VLAN's name or VLAN ID (VID) number. For example, if a VLAN configured in the switch has a VID of 100 and is named **vlan100**, you could configure the RADIUS server to use either "100" or "vlan100" to specify the VLAN.

After the RADIUS server validates a client's username and password, the RADIUS server returns an Access-Accept packet that contains the VLAN assignment and the following attributes for use in the authentication session:

- **Egress-VLANID:** Configures an optional, egress VLAN ID for either tagged or untagged packets (RFC 4675).
- **Egress-VLAN-Name:** Configures an optional, egress VLAN for either tagged or untagged packets when the VLAN ID is not known (RFC 4675).
- **Tunnel-Type, Tunnel-Medium-Type, and Tunnel-Private-Group-ID:** Tunnel attributes that specify an untagged VLAN assignment (RFC 3580).

Tunnel (untagged VLAN) attributes may be included in the same RADIUS packet as the Egress-VLANID and Egress-VLAN-Name attributes. These attributes are not mutually exclusive.

The switch processes the VLAN information returned from the remote RADIUS server for each successfully 802.1X-, Web-, and MAC-authenticated client (user). The VLAN information is part of the user's profile stored in the RADIUS server's database and is applied if the VLANs exist on the switch.

The support for RADIUS-assigned tagged and untagged VLAN configuration on an authenticated port allows you to use IDM to dynamically configure tagged and untagged VLANs as required for different client devices, such as PCs and IP phones, that share the same switch port.

Additional RADIUS Attributes

The following attributes are included in Access-Request and Access-Accounting packets sent from the switch to the RADIUS server to advertise switch capabilities, report information on authentication sessions, and dynamically reconfigure authentication parameters:

- **MS-RAS-Vendor (RFC 2548):** Allows ProCurve switches to inform a Microsoft RADIUS server that the switches are from ProCurve Networking. This feature assists the RADIUS server in its network configuration.
- **HP-capability-advert:** A ProCurve proprietary RADIUS attribute that allows a switch to advertise its current capabilities to the RADIUS server for port-based (MAC, Web, or 802.1X) authentication; for example, HP VSAs for port QoS, ingress rate-limiting, IDM filter rules, RFC 4675 QoS and VLAN attributes, and RFC 3580 VLAN-related attributes.

The RADIUS server uses this information to make a more intelligent policy decision on the configuration settings to return to the switch for a client session.

- **HP-acct-terminate-cause:** A ProCurve proprietary RADIUS accounting attribute that allows a switch to report to the RADIUS server why an authentication session was terminated. This information allows customers to diagnose network operational problems and generate reports on terminated sessions. This attribute provides extended information on the statistics provided by the acct-terminate-cause attribute.
- **change-of-authorization (RFC 3576: Dynamic Authorization Extensions to RADIUS):** A mechanism that allows a RADIUS server to dynamically terminate or change the authorization parameters (such as VLAN assignment) used in an active client session on the switch. The switch (NAS) does not have to initiate the exchange.

For example, for security reasons you may want to limit the network services granted to an authenticated user. In this case, you can change the user profile on the RADIUS server and have the new authorization settings take effect immediately in the active client session. The change-of-authorization attribute provides the mechanism to dynamically update an active client session with a new user policy that is sent in RADIUS packets.

Configuring RADIUS Accounting

RADIUS Accounting Commands	Page
[no] radius-server host < ip-address >	5-40
[acct-port < port-number >]	5-40
[key < key-string >]	5-40
[no] aaa accounting < exec network system commands > < start-stop stop-only > radius	5-44
[no] aaa accounting update periodic < 1 - 525600 > (in minutes)	5-44
[no] aaa accounting suppress null-username	5-44
show accounting	5-49
show accounting sessions	5-50
show radius accounting	5-50

Note

This section assumes you have already:

- Configured RADIUS authentication on the switch for one or more access methods
- Configured one or more RADIUS servers to support the switch

If you have not already done so, refer to “General RADIUS Setup Procedure” on page 5-7 before continuing here.

RADIUS accounting collects data about user activity and system events and sends it to a RADIUS server when specified events occur on the switch, such as a logoff or a reboot. The switches covered in this guide support four types of accounting services:

- **Network accounting:** Provides records containing the information listed below on clients directly connected to the switch and operating under Port-Based Access Control (802.1X):
 - Acct-Authentic
 - Acct-Delay-Time
 - Acct-Input-Octets
 - Acct-Input-Packets
 - Acct-Output-Octets
 - Acct-Output-Packets
 - Acct-Session-Id
 - Acct-Session-Time
 - Acct-Status-Type
 - Acct-Terminate-Cause
 - Called-Station-Id
 - HP-acct-terminate-cause
 - MS-RAS-Vendor
 - NAS-Identifier
 - NAS-IP-Address
 - NAS-Port
 - Service-Type
 - Username

- **Exec accounting:** Provides records holding the information listed below about login sessions (console, Telnet, and SSH) on the switch:

- | | | |
|---------------------|------------------------|------------------|
| • Acct-Authentic | • Acct-Status-Type | • NAS-Identifier |
| • Acct-Delay-Time | • Acct-Terminate-Cause | • NAS-IP-Address |
| • Acct-Session-Id | • Calling-Station-Id | • Service-Type |
| • Acct-Session-Time | • MS-RAS-Vendor | • Username |

- **System accounting:** Provides records containing the information listed below when system events occur on the switch, including system reset, system boot, and enabling or disabling of system accounting.

- | | | |
|---------------------|------------------------|------------------|
| • Acct-Authentic | • Acct-Terminate-Cause | • NAS-IP-Address |
| • Acct-Delay-Time | • Calling-Station-Id | • Service-Type |
| • Acct-Session-Id | • MS-RAS-Vendor | • Username |
| • Acct-Session-Time | • NAS-Identifier | |

- **Commands accounting:** Provides records containing information after the execution of a command.

- **RADIUS accounting with IP attribute:** The RADIUS Attribute 8 (Framed-IP-Address) feature provides the RADIUS server with information about the client's IP address after the client is authenticated. DHCP snooping is queried for the IP address of the client, so DHCP snooping must be enabled for the VLAN of which the client is a member.

When the switch begins communications with the RADIUS server it sends the IP address of the client requesting access to the RADIUS server as RADIUS Attribute 8 (Framed-IP-Address) in the RADIUS accounting request. The RADIUS server can use this information to build a map of usernames and addresses.

It may take a minute or longer for the switch to learn the IP address and then send the accounting packet with the Framed-IP-Address attribute to the RADIUS server. If the switch does not learn the IP address after a minute, it sends the accounting request packet to the RADIUS server without the Framed-IP-Address attribute. If the IP address is learned at a later time, it will be included in the next accounting request packet sent.

The switch forwards the accounting information it collects to the designated RADIUS server, where the information is formatted, stored, and managed by the server. For more information on this aspect of RADIUS accounting, refer to the documentation provided with your RADIUS server.

Operating Rules for RADIUS Accounting

- You can configure up to four types of accounting to run simultaneously: exec, system, network, and commands.
- RADIUS servers used for accounting are also used for authentication.
- The switch must be configured to access at least one RADIUS server.
- RADIUS servers are accessed in the order in which their IP addresses were configured in the switch. Use **show radius** to view the order. As long as the first server is accessible and responding to authentication requests from the switch, a second or third server will not be accessed. (For more on this topic, refer to “Changing RADIUS-Server Access Order” on page 5-50.)
- If access to a RADIUS server fails during a session, but after the client has been authenticated, the switch continues to assume the server is available to receive accounting data. Thus, if server access fails during a session, it will not receive accounting data transmitted from the switch.

Steps for Configuring RADIUS Accounting

1. Configure the switch for accessing a RADIUS server.

You can configure a list of up to three RADIUS servers (one primary, two backup). The switch operates on the assumption that a server can operate in both accounting and authentication mode. (Refer to the documentation for your RADIUS server application.)

- Use the same **radius-server host** command that you would use to configure RADIUS authentication. Refer to “3. Configure the Switch To Access a RADIUS Server” on page 5-14.
- Provide the following:
 - A RADIUS server IP address.
 - Optional—a UDP destination port for authentication requests. Otherwise the switch assigns the default UDP port (1812; recommended).
 - Optional—if you are also configuring the switch for RADIUS authentication, and need a unique encryption key for use during authentication sessions with the RADIUS server you are designating, configure a server-specific key. This key overrides the global encryption key you can also configure on the switch, and

must match the encryption key used on the specified RADIUS server. For more information, refer to the “[**key < key-string >**]” parameter on page 5-14. (Default: null)

2. Configure accounting types and the controls for sending reports to the RADIUS server.
 - **Accounting types:** exec (page 5-38), network (page 5-37), commands (page 5-38), or system (page 5-38)
 - **Trigger for sending accounting reports to a RADIUS server:** At session start and stop or only at session stop
3. (Optional) Configure session blocking and interim updating options
 - **Updating:** Periodically update the accounting data for sessions-in-progress
 - **Suppress accounting:** Block the accounting session for any unknown user with no username access to the switch

1. Configure the Switch To Access a RADIUS Server

Before you configure the actual accounting parameters, you should first configure the switch to use a RADIUS server. This is the same as the process described on page 5-14. You need to repeat this step here only if you have not yet configured the switch to use a RADIUS server, your server data has changed, or you need to specify a non-default UDP destination port for accounting requests. Note that switch operation expects a RADIUS server to accommodate both authentication and accounting.

Syntax: [no] radius-server host < ip-address > [oobm]

*Adds a server to the RADIUS configuration or (with **no**) deletes a server from the configuration.*

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the RADIUS traffic will go through the out-of-band management (OOBM) port.*

[acct-port < port-number >]

Optional. Changes the UDP destination port for accounting requests to the specified RADIUS server. If you do not use this option, the switch automatically assigns the default accounting port number. (Default: 1813)

[key < key-string >]

Optional. Specifies an encryption key for use during accounting or authentication sessions with the specified server. This key must match the encryption key used on the RADIUS server. Use this command only if the specified server requires a different encryption key than configured for the global encryption key.

Note: When you save the config file using Xmodem or TFTP, the key information is not saved in the file. This causes Radius authentication to fail when the config file is loaded back onto the switch.

(For a more complete description of the **radius-server** command and its options, turn to page 5-14.)

For example, suppose you want the switch to use the RADIUS server described below for both authentication and accounting purposes.

- IP address: 10.33.18.151
- A non-default UDP port number of 1750 for accounting.

For this example, assume that all other RADIUS authentication parameters for accessing this server are acceptable at their default settings, and that RADIUS is already configured as an authentication method for one or more types of access to the switch (Telnet, Console, etc.).

```
ProCurve(config)# radius-server host 10.33.18.151 acct-port 1750 key source0151
ProCurve(config)# write mem
ProCurve(config)# show radius
```

Status and Counters - General RADIUS Information

Deadtime(min) : 5
Timeout(secs) : 3
Retransmit Attempts : 2
Global Encryption Key :

Server IP Addr	Auth Port	Acct Port	Encryption Key
10.33.18.151	1812	1750	source0151

Because the radius-server command includes an **acct-port** element with a non-default 1750, the switch assigns this value to the accounting port UDP port numbers. Because auth-port was not included in the command, the authentication UDP port is set to the default 1812.

Figure 5-11. Example of Configuring for a RADIUS Server with a Non-Default Accounting UDP Port Number

The radius-server command as shown in figure 5-11, above, configures the switch to use a RADIUS server at IP address 10.33.18.151, with a (non-default) UDP accounting port of 1750, and a server-specific key of “source0151”.

2. Configure Accounting Types and the Controls for Sending Reports to the RADIUS Server

Select the Accounting Type(s):

- **Exec:** Use **exec** if you want to collect accounting information on login sessions on the switch via the console, Telnet, or SSH. (See also “Accounting Services” on page 5-4.)
- **System:** Use **system** if you want to collect accounting data when:
 - A system boot or reload occurs
 - System accounting is turned on or off

Note that there is no time span associated with using the **system** option. It simply causes the switch to transmit whatever accounting data it currently has when one of the above events occurs.

- **Network:** Use Network if you want to collect accounting information on 802.1X port-based-access users connected to the physical ports on the switch to access the network. (See also “Accounting Services” on page 4.)
- **Commands:** When commands authorization is enabled, a record accounting notice is sent after the execution of a command.
- **Web or MAC:** You can also use Web or MAC to collect accounting information.

Determine how you want the switch to send accounting data to a RADIUS server:

- **Start-Stop:**
 - Send a start record accounting notice at the beginning of the accounting session and a stop record notice at the end of the session. Both notices include the latest data the switch has collected for the requested accounting type (Network, Exec, Commands, or System).
 - Do not wait for an acknowledgement.

The system option (page 5-42) ignores **start-stop** because the switch sends the accumulated data only when there is a reboot, reload, or accounting on/off event.

- **Stop-Only:**
 - Send a stop record accounting notice at the end of the accounting session. The notice includes the latest data the switch has collected for the requested accounting type (Network, Exec, Commands, or System).
 - Do not wait for an acknowledgment.

The system option (page 5-42) always delivers **stop-only** operation because the switch sends the accumulated data only when there is a reboot, reload, or accounting on/off event.

Syntax: [no] aaa accounting < exec | network | system | commands > < start-stop | stop-only > radius

Configures RADIUS accounting type and how data will be sent to the RADIUS server.

For example, to configure RADIUS accounting on the switch with **start-stop** for exec functions and **stop-only** for system functions:

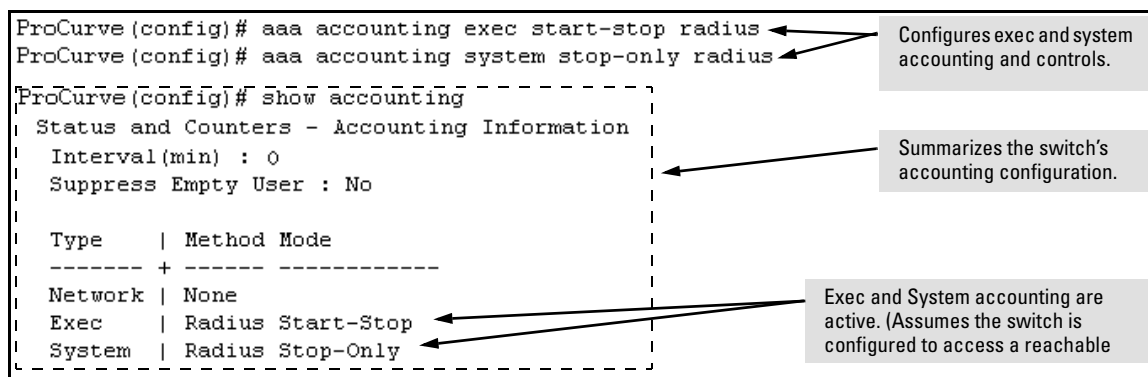


Figure 5-12. Example of Configuring Accounting Types

3. (Optional) Configure Session Blocking and Interim Updating Options

These optional parameters give you additional control over accounting data.

- **Updates:** In addition to using a Start-Stop or Stop-Only trigger, you can optionally configure the switch to send periodic accounting record updates to a RADIUS server.
- **Suppress:** The switch can suppress accounting for an unknown user having no username.

Syntax: [no] aaa accounting update periodic < 1 - 525600>

*Sets the accounting update period for all accounting sessions on the switch. (The **no** form disables the update function and resets the value to zero.) (Default: zero; disabled).*

Syntax: [no] aaa accounting suppress null-username

Disables accounting for unknown users having no username. (Default: suppression disabled)

To continue the example in figure 5-12, suppose that you wanted the switch to:

- Send updates every 10 minutes on in-progress accounting sessions.
- Block accounting for unknown users (no username).

```
ProCurve(config)# aaa accounting update periodic 10
ProCurve(config)# aaa accounting suppress null-username

ProCurve(config)# show accounting
Status and Counters - Accounting Information
Interval(min) : 10
Suppress Empty User : Yes

Type      | Method Mode
-----+-----
Network   | None
Exec      | Radius Start-Stop
System    | Radius Stop-Only
```




Figure 5-13. Example of Optional Accounting Update Period and Accounting Suppression on Unknown User

Viewing RADIUS Statistics

General RADIUS Statistics

Syntax: show radius [host < ip-addr >]

*Shows general RADIUS configuration, including the server IP addresses. Optional form shows data for a specific RADIUS host. To use **show radius**, the server's IP address must be configured in the switch, which. requires prior use of the **radius-server host** command. (See "Configuring RADIUS Accounting" on page 5-37.)*

```
ProCurve(config)# show radius
Status and Counters - General RADIUS Information
  Deadttime(min) : 5
  Timeout(secs) : 10
  Retransmit Attempts : 2
  Global Encryption Key : myg10balkey

  Server IP Addr  Auth  Acct
                  Port  Port  Encryption Key
  -----
  192.33.12.65    1812  1813  my65key
```

Figure 5-14. Example of General RADIUS Information from Show Radius Command

```
ProCurve(config)# show radius host 192.33.12.65
Status and Counters - RADIUS Server Information

Server IP Addr : 192.33.12.65
Authentication UDP Port : 1812           Accounting UDP Port : 1813
Round Trip Time      : 2                 Round Trip Time      : 7
Pending Requests     : 0                 Pending Requests     : 0
Retransmissions      : 0                 Retransmissions      : 0
Timeouts             : 0                 Timeouts             : 0
Malformed Responses  : 0                 Malformed Responses  : 0
Bad Authenticators   : 0                 Bad Authenticators   : 0
Unknown Types        : 0                 Unknown Types        : 0
Packets Dropped      : 0                 Packets Dropped      : 0
Access Requests      : 2                 Accounting Requests   : 2
Access Challenges     : 0                 Accounting Responses  : 2
Access Accepts        : 2
Access Rejects        : 0
```

Figure 5-15. RADIUS Server Information From the Show Radius Host Command

Term	Definition
Round Trip Time	The time interval between the most recent Accounting-Response and the Accounting-Request that matched it from this RADIUS accounting server.
PendingRequests	The number of RADIUS Accounting-Request packets sent to this server that have not yet timed out or received a response. This variable is incremented when an accounting-Request is sent and decremented due to receipt of an Accounting-Response, a timeout or a retransmission.
Retransmissions	The number of RADIUS Accounting-Request packets retransmitted to this RADIUS accounting server. Retransmissions include retries where the Identifier and Acct-Delay have been updated, as well as those in which they remain the same.
Timeouts	The number of accounting timeouts to this server. After a timeout the client may retry to the same server, send to a different server, or give up. A retry to the same server is counted as a retransmit as well as a timeout. A send to a different server is counted as an Accounting-Request as well as a timeout.
Malformed Responses	The number of malformed RADIUS Accounting-Response packets received from this server. Malformed packets include packets with an invalid length. Bad authenticators and unknown types are not included as malformed accounting responses.
Bad Authenticators	The number of RADIUS Accounting-Response packets which contained invalid authenticators received from this server.
Unknown Types	The number of RADIUS packets of unknown type which were received from this server on the accounting port.
Packets Dropped	The number of RADIUS packets which were received from this server on the accounting port and dropped for some other reason.

Term	Definition
Requests	The number of RADIUS Accounting-Request packets sent. This does not include retransmissions.
AccessChallenges	The number of RADIUS Access-Challenge packets (valid or invalid) received from this server.
AccessAccepts	The number of RADIUS Access-Accept packets (valid or invalid) received from this server.
AccessRejects	The number of RADIUS Access-Reject packets (valid or invalid) received from this server.
Responses	The number of RADIUS packets received on the accounting port from this server.

RADIUS Authentication Statistics

Syntax: show authentication

Displays the primary and secondary authentication methods configured for the Console, Telnet, Port-Access (802.1X), and SSH methods of accessing the switch. Also displays the number of access attempts currently allowed in a session.

show radius authentication

Displays NAS identifier and data on the configured RADIUS server and the switch's interactions with this server.

*(Requires prior use of the **radius-server host** command to configure a RADIUS server IP address in the switch. See "Configuring RADIUS Accounting" on page 5-37.)*

```
ProCurve(config)# show authentication
```

Status and Counters - Authentication Information				
Login Attempts : 2				
Respect Privilege : Disabled				
Access Task	Login Primary	Login Secondary	Enable Primary	Enable Secondary
Console	Local	None	Local	None
Telnet	Radius	None	Radius	None
Port-Access	Local			
Webui	Local	None	Local	None
SSH	Radius	None	Radius	None
Web-Auth	ChapRadius			
MAC-Auth	ChapRadius			

Figure 5-16. Example of Login Attempt and Primary/Secondary Authentication Information from the Show Authentication Command

```
ProCurve(config)# show radius authentication
Status and Counters - RADIUS Authentication Information

NAS Identifier : ProCurve.
Invalid Server Addresses : 0

      UDP
Server IP Addr  Port  Timeouts  Requests  Challenges  Accepts  Rejects
-----
192.33.12.65   1812  0         2          0           2        0
```

Figure 5-17. Example of RADIUS Authentication Information from a Specific Server

RADIUS Accounting Statistics

Syntax: show accounting

Lists configured accounting interval, “Empty User” suppression status, accounting types, methods, and modes.

show radius accounting

*Lists accounting statistics for the RADIUS server(s) configured in the switch (using the **radius-server host** command).*

show accounting sessions

Lists the accounting sessions currently active on the switch.

```
HPswitch # show accounting

Status and Counters - Accounting Information

Interval(min) : 5
Suppress Empty User : Yes

Type      | Method Mode
-----+-----
Network  | None
Exec     | Radius Start-Stop
System   | Radius Stop-Only
```

Figure 5-18. Listing the Accounting Configuration in the Switch


```
ProCurve # show radius accounting
Status and Counters - RADIUS Accounting Information
NAS Identifier : HPswitch
Invalid Server Addresses : 0

      UDP
Server IP Addr  Port  Timeouts  Requests  Responses
-----
192.33.12.65    1813  0         1          1
```

Figure 5-19. Example of RADIUS Accounting Information for a Specific Server

```
ProCurve # show accounting sessions

Active Accounted actions on CONSOLE, User radius Priv 2,
Session ID 1, EXEC Accounting record, 00:02:32 Elapsed
```

Figure 5-20. Example Listing of Active RADIUS Accounting Sessions on the Switch

Changing RADIUS-Server Access Order

The switch tries to access RADIUS servers according to the order in which their IP addresses are listed by the **show radius** command. Also, *when you add a new server IP address, it is placed in the highest empty position in the list.*

Adding or deleting a RADIUS server IP address leaves an empty position, but does not change the position of any other server addresses in the list. For example if you initially configure three server addresses, they are listed in the order in which you entered them. However, if you subsequently remove the second server address in the list and add a new server address, the new address will be placed second in the list.

Thus, to move a server address up in the list, you must delete it from the list, ensure that the position to which you want to move it is vacant, and then re-enter it. For example, suppose you have already configured the following three RADIUS server IP addresses in the switch:

```
ProCurve # show radius
Status and Counters - General RADIUS Information
  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key :

  Auth  Acct
Server IP Addr  Port  Port  Encryption Key
-----
10.10.10.1      1812  1813
10.10.10.2      1812  1813
10.10.10.3      1812  1813
```

RADIUS server IP addresses listed in the order in which the switch will try to access them. In this case, the server at IP address 1.1.1.1 is first.

Note: If the switch successfully accesses the first server, it does not try to access any other servers in the list, even if the client is denied access by the first server.

Figure 5-21. Search Order for Accessing a RADIUS Server

To exchange the positions of the addresses so that the server at 10.10.10.003 will be the first choice and the server at 10.10.10.001 will be the last, you would do the following:

1. Delete 10.10.10.003 from the list. This opens the third (lowest) position in the list.
2. Delete 10.10.10.001 from the list. This opens the first (highest) position in the list.
3. Re-enter 10.10.10.003. Because the switch places a newly entered address in the highest-available position, this address becomes first in the list.
4. Re-enter 10.10.10.001. Because the only position open is the third position, this address becomes last in the list.

RADIUS Authentication, Authorization, and Accounting

Changing RADIUS-Server Access Order

```
ProCurve(config)# no radius host 10.10.10.003
ProCurve(config)# no radius host 10.10.10.001
ProCurve(config)# radius host 10.10.10.003
ProCurve(config)# radius host 10.10.10.001

ProCurve(config)# show radius
```

Status and Counters - General RADIUS Information

```
Deadtime(min) : 0
Timeout(secs) : 5
Retransmit Attempts : 3
Global Encryption Key :
```

Server IP Addr	Auth Port	Acct Port	Encryption Key
10.10.10.3	1812	1813	
10.10.10.2	1812	1813	
10.10.10.1	1812	1813	

Removes the "003" and "001" addresses from the RADIUS server list.

Inserts the "003" address in the first position in the RADIUS server list, and inserts the "001" address in the last position in the list.

Shows the new order in which the switch searches for a RADIUS server.

Figure 5-22. Example of New RADIUS Server Search Order

Messages Related to RADIUS Operation

Message	Meaning
Can't reach RADIUS server < x.x.x.x >.	A designated RADIUS server is not responding to an authentication request. Try pinging the server to determine whether it is accessible to the switch. If the server is accessible, then verify that the switch is using the correct encryption key and that the server is correctly configured to receive an authentication request from the switch.
No server(s) responding.	The switch is configured for and attempting RADIUS authentication, however it is not receiving a response from a RADIUS server. Ensure that the switch is configured to access at least one RADIUS server. (Use show radius .) If you also see the message Can't reach RADIUS server < x.x.x.x >, try the suggestions listed for that message.
Not legal combination of authentication methods.	Indicates an attempt to configure local as both the primary and secondary authentication methods. If local is the primary method, then none must be the secondary method.

Configuring Secure Shell (SSH)

Contents

Overview	6-2
Terminology	6-3
Prerequisite for Using SSH	6-5
Public Key Formats	6-5
Steps for Configuring and Using SSH for Switch and Client Authentication	6-6
General Operating Rules and Notes	6-8
Configuring the Switch for SSH Operation	6-9
1. Assigning a Local Login (Operator) and Enable (Manager) Password	6-9
2. Generating the Switch's Public and Private Key Pair	6-10
Configuring Key Lengths	6-13
3. Providing the Switch's Public Key to Clients	6-13
4. Enabling SSH on the Switch and Anticipating SSH Client Contact Behavior	6-16
5. Configuring the Switch for SSH Authentication	6-21
6. Use an SSH Client To Access the Switch	6-24
Further Information on SSH Client Public-Key Authentication	6-25
Messages Related to SSH Operation	6-30
Logging Messages	6-31
Debug Logging	6-32

Overview

Feature	Default	Menu	CLI	Web
Generating a public/private key pair on the switch	No	n/a	page 6-10	n/a
Using the switch's public key	n/a	n/a	page 6-13	n/a
Enabling SSH	Disabled	n/a	page 6-15	n/a
Enabling client public-key authentication	Disabled	n/a	pages 6-21, 6-24	n/a
Enabling user authentication	Disabled	n/a	page 6-20	n/a

The switches covered in this guide use Secure Shell version 2 (SSHv2) to provide remote access to management functions on the switches via encrypted paths between the switch and management station clients capable of SSH operation.

SSH provides Telnet-like functions but, unlike Telnet, SSH provides encrypted, authenticated transactions. The authentication types include:

- Client public-key authentication
- Switch SSH and user password authentication

Client Public Key Authentication (Login/Operator Level) with User Password Authentication (Enable/Manager Level). This option uses one or more public keys (from clients) that must be stored on the switch. Only a client with a private key that matches a stored public key can gain access to the switch. (The same private key can be stored on one or more clients.)

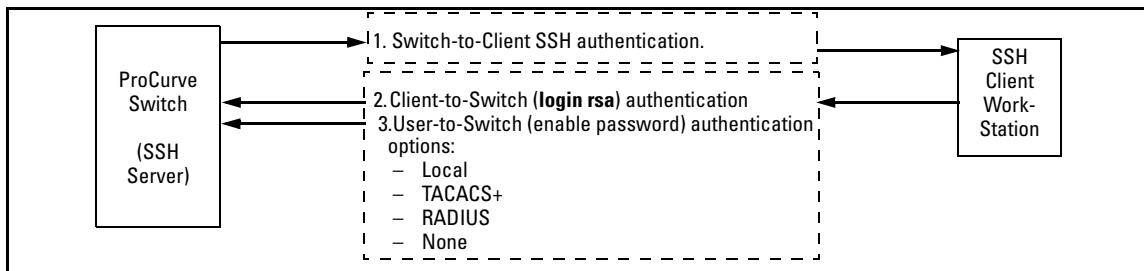


Figure 6-1. Client Public Key Authentication Model

Note

SSH in ProCurve switches is based on the OpenSSH software toolkit. For more information on OpenSSH, visit www.openssh.com.

Switch SSH and User Password Authentication . This option is a subset of the client public-key authentication shown in figure 6-1. It occurs if the switch has SSH enabled but does not have login access (**login public-key**) configured to authenticate the client's key. As in figure 6-1, the switch authenticates itself to SSH clients. Users on SSH clients then authenticate themselves to the switch (login and/or enable levels) by providing passwords stored locally on the switch or on a TACACS+ or RADIUS server. However, the client does not use a key to authenticate itself to the switch.

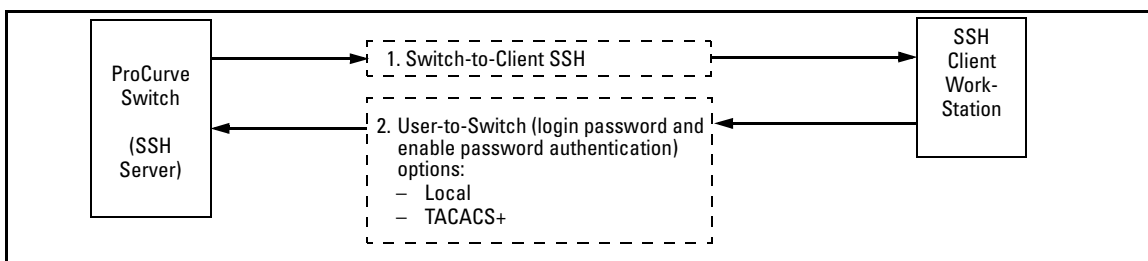


Figure 6-2. Switch/User Authentication

Terminology

- **SSH Server:** A ProCurve switch with SSH enabled.
- **Key Pair:** A pair of keys generated by the switch or an SSH client application. Each pair includes a public key, that can be read by anyone and a private key held internally in the switch or by a client.
- **PEM (Privacy Enhanced Mode):** Refers to an ASCII-formatted client public-key that has been encoded for portability and efficiency. SSHv2 client public-keys are typically stored in the PEM format. See figure 6-3 for an example of PEM-encoded ASCII keys.
- **Private Key:** An internally generated key used in the authentication process. A private key generated by the switch is not accessible for viewing or copying. A private key generated by an SSH client application is typically stored in a file on the client device and, together with its public key counterpart, can be copied and stored on multiple devices.
- **Public Key:** An internally generated counterpart to a private key. A device's public key is used to authenticate the device to other devices.

- **Enable Level:** Manager privileges on the switch.
- **Login Level:** Operator privileges on the switch.
- **Local password or username:** A Manager-level or Operator-level password configured in the switch.
- **SSH Enabled:** (1) A public/private key pair has been generated on the switch (**generate ssh [dsa | rsa]**) and (2) SSH is enabled (**ip ssh**). (You can generate a key pair without enabling SSH, but you cannot enable SSH without first generating a key pair. See “2. Generating the Switch’s Public and Private Key Pair” on page 6-10 and “4. Enabling SSH on the Switch and Anticipating SSH Client Contact Behavior” on page 6-15.)

Prerequisite for Using SSH

Before using the switch as an SSH server, you must install a publicly or commercially available SSH client application on the computer(s) you use for management access to the switch. If you want client public-key authentication (page 6-2), then the client program must have the capability to generate or import keys.

Public Key Formats

Any client application you use for client public-key authentication with the switch must have the capability to export public keys. The switch can accept keys in the PEM-Encoded ASCII Format or in the Non-Encoded ASCII format.

The diagram shows a single line of text representing a PEM-Encoded public key. The text is enclosed in double quotes. The first part of the text is a comment: "Pub Key Gen 21 Dec 2001 12:01". This is followed by a space, then the actual public key data: "A1B3Nz1y2+orEhYL . . . Q8D8qDH1ozu1c". The string ends with "=== End of Pub Key ===". Two annotations are present: a grey box labeled "Comment describing public" with an arrow pointing to the first part of the string, and another grey box labeled "Beginning of actual SSHv2 public key in PEM-Encoded" with an arrow pointing to the start of the key data.

```
"Pub Key Gen 21 Dec 2001 12:01" A1B3Nz1y2+orEhYL . . . Q8D8qDH1ozu1c"=== End of Pub Key ==="
```

Figure 6-3. Example of Public Key in PEM-Encoded ASCII Format Common for SSHv2 Clients

Steps for Configuring and Using SSH for Switch and Client Authentication

For two-way authentication between the switch and an SSH client, you must use the login (Operator) level.

Table 6-1. SSH Options

Switch Access Level	Primary SSH Authentication	Authenticate Switch Public Key to SSH Clients?	Authenticate Client Public Key to the Switch?	Primary Switch Password Authentication	Secondary Switch Password Authentication
Operator (Login) Level	ssh login rsa	Yes	Yes ¹	No ¹	local or none
	ssh login Local	Yes	No	Yes	none
	ssh login TACACS	Yes	No	Yes	local or none
	ssh login RADIUS	Yes	No	Yes	local or none
Manager (Enable) Level	ssh enable local	Yes	No	Yes	none
	ssh enable tacacs	Yes	No	Yes	local or none
	ssh enable radius	Yes	No	Yes	local or none

¹ For **ssh login public-key**, the switch uses client public-key authentication instead of the switch password options for primary authentication.

The general steps for configuring SSH include:

A. Client Preparation

1. Install an SSH client application on a management station you want to use for access to the switch. (Refer to the documentation provided with your SSH client application.)
2. Optional—If you want the switch to authenticate a client public-key on the client:
 - a. Either generate a public/private key pair on the client computer (if your client application allows) or import a client key pair that you have generated using another SSH application.
 - b. Copy the client public key into an ASCII file on a TFTP server accessible to the switch and download the client public key file to the switch. (The client public key file can hold up to 10 client keys.) This topic is covered under “To Create a Client-Public-Key Text File” on page 6-26.

B. Switch Preparation

1. Assign a login (Operator) and enable (Manager) password on the switch (page 6-10).
2. Generate a public/private key pair on the switch (page 6-10).

You need to do this only once. The key remains in the switch even if you reset the switch to its factory-default configuration. (You can remove or replace this key pair, if necessary.)
3. Copy the switch's public key to the SSH clients you want to access the switch (page 6-13).
4. Enable SSH on the switch (page 6-15).
5. Configure the primary and secondary authentication methods you want the switch to use. In all cases, the switch will use its host-public-key to authenticate itself when initiating an SSH session with a client.
 - SSH Login (Operator) options:
 - Option A:

Primary: Local, TACACS+, or RADIUS password
Secondary: Local password or none. If the primary option is local, the secondary option must be none.
 - Option B:

Primary: Client public-key authentication (**login public-key** — page 6-24)
Secondary: none

Note that if you want the switch to perform client public-key authentication, you must configure the switch with Option B.
- SSH Enable (Manager) options:

Primary: Local, TACACS+, or RADIUS
Secondary: Local password or none. If the primary option is local, the secondary option must be none.
6. Use your SSH client to access the switch using the switch's IP address or DNS name (if allowed by your SSH client application). Refer to the documentation provided with the client application.

General Operating Rules and Notes

- Public keys generated on an SSH client must be exportable to the switch. The switch can only store 10 client key pairs.
- The switch's own public/private key pair and the (optional) client public key file are stored in the switch's flash memory and are not affected by reboots or the **erase startup-config** command.
- Once you generate a key pair on the switch you should avoid re-generating the key pair without a compelling reason. Otherwise, you will have to re-introduce the switch's public key on all management stations (clients) you previously set up for SSH access to the switch. In some situations this can temporarily allow security breaches.
- The switch does not support outbound SSH sessions. Thus, if you Telnet from an SSH-secure switch to another SSH-secure switch, *the session is not secure*.
- With SSH running, the switch allows one console session and up to five other sessions (SSH and/or Telnet).

Configuring the Switch for SSH Operation

SSH-Related Commands in This Section	Page
show ip ssh	6-19
show crypto client-public-key [<manager operator>] [keylist-str] [< babble fingerprint>]	6-27
show crypto host-public-key [< babble fingerprint >]	6-14
show authentication	6-23
crypto key < generate zeroize > [autorun-key [rsa] cert [rsa] <keysize> ssh [dsa rsa [bits <keysize>]]	6-11
ip ssh	6-16
cipher <cipher-type>	6-17
filetransfer	6-17
ip-version	6-17
mac	6-18
port < 1 - 65535 default >	6-17
timeout < 5 - 120 >	6-17
listen <oobm data both>	6-18
aaa authentication ssh	
login < local tacacs radius public-key >	6-20, 6-22
< local none >	6-20
enable < tacacs radius local >	6-20
< local none >	6-20
copy tftp pub-key-file <tftp server IP> <public key file> [<append manager operator>] [oobm]	6-27
clear crypto client-public-key [keylist-str]	6-29

1. Assigning a Local Login (Operator) and Enable (Manager) Password

At a minimum, ProCurve recommends that you always assign at least a Manager password to the switch. Otherwise, under some circumstances, anyone with Telnet, web, or serial port access could modify the switch's configuration.

To Configure Local Passwords. You can configure both the Operator and Manager password with one command.

Syntax: password < manager | operator | all >

```
ProCurve(config)# password all
New password for Operator: *****
Please retype new password for Operator: *****
New password for Manager: *****
Please retype new password for Manager: *****
ProCurve(config)#
```

Figure 6-4. Example of Configuring Local Passwords

2. Generating the Switch's Public and Private Key Pair

You must generate a public and private host key pair on the switch. The switch uses this key pair, along with a dynamically generated session key pair to negotiate an encryption method and session with an SSH client trying to connect to the switch.

The host key pair is stored in the switch's flash memory, and only the public key in this pair is readable. The public key should be added to a "known hosts" file (for example, `$HOME/.ssh/known_hosts` on UNIX systems) on the SSH clients which should have access to the switch. Some SSH client applications automatically add the switch's public key to a "known hosts" file. Other SSH applications require you to manually create a known hosts file and place the switch's public key in the file. (Refer to the documentation for your SSH client application.)

(The session key pair mentioned above is not visible on the switch. It is a temporary, internally generated pair used for a particular switch/client session, and then discarded.)

Note

When you generate a host key pair on the switch, the switch places the key pair in flash memory (and not in the running-config file). Also, the switch maintains the key pair across reboots, including power cycles. You should consider this key pair to be “permanent”; that is, avoid re-generating the key pair without a compelling reason. Otherwise, you will have to re-introduce the switch’s public key on all management stations you have set up for SSH access to the switch using the earlier pair.

Removing (zeroing) the switch’s public/private key pair renders the switch unable to engage in SSH operation and automatically disables IP SSH on the switch. (To verify whether SSH is enabled, execute **show ip ssh**.) However, any active SSH sessions will continue to run, unless explicitly terminated with the CLI 'kill' command.

To Generate or Erase the Switch’s Public/Private Host Key Pair.

Because the host key pair is stored in flash instead of the running-config file, it is not necessary to use **write memory** to save the key pair. Erasing the key pair automatically disables SSH.

Syntax: `crypto key generate <autorun-key [rsa] | cert [rsa] <keysize> | ssh [dsa | rsa] bits <keysize>>`

Installs authentication files for ssh or https server, or for autorun.

`autorun-key`

Install RSA key for autorun. See “Configuring Autorun on the Switch” in Appendix A of the Management and Configuration Guide for more information.

`cert`

Install RSA key for https certificate. See “Configuring the Switch for SSL Operation” on page 7-7 in this guide for more information.

`ssh [dsa | rsa]`

Install host key for ssh server. Specify the key type as DSA or RSA.

`bits <keysize>`

Specify the key size (in bits). See Table 6-2.

`zeroize <ssh | cert | autorun [rsa]>`

Erases the switch’s public/private key pair and disables SSH operation.

show crypto host-public-key

Displays switch's public key. Displays the version 1 and version 2 views of the key.

[babble]

Displays hashes of the switch's public key in phonetic format. (See "Displaying the Public Key" on page 6-14.)

[fingerprint]

Displays fingerprints of the switch's public key in hexadecimal format. (See "Displaying the Public Key" on page 6-14.)

For example, to generate and display a new key:

```
ProCurve(config)# crypto key generate ssh rsa
Installing new RSA key.  If the key/entropy cache is
depleted, this could take up to a minute.
ProCurve(config)# show crypto host-public-key

-----
SSH host public key file
Version 1 format:
-----
896 35 3219295003103011452137203169501232714847265325085720757925409572738582167
49173126937413223781326827636154399173519641900117298772018339016754333892248319
41759125186557710233731689070801858880718460531164552600040416069890120011153581
9449254242176260739141950918771764467
-----
Version 2 format:
-----
ssh-rsa AAAAB3NzaClyc2EAAAABIAAAHEAnAAApdhq13Jynrs7j4lDUm8ivVm8ld2mZU5e+YZWp/T6
QzP2RsDDMZLbAHHIBrxPLjW/bRogpYD01WuV0hTojEVjqeVuXbwmdDny0gBc06olePwdrbQ+FZevERiA
JYG3C8NCzCRD/djXeI7FmRps8w==
-----
```

Host Public Key for the Switch

Version 1 and Version 2 views of same host public key

Figure 6-5. Example of Generating a Public/Private Host Key Pair for the Switch

The 'show crypto host-public-key' displays it in two different formats because your client may store it in either of these formats after learning the key. If you wish to compare the switch key to the key as stored in your client's known-hosts file, note that the formatting and comments need not match. For version 1 keys, the three numeric values bit size, exponent <e>, and modulus <n> must match; for PEM keys, only the PEM-encoded string itself must match.

Notes

"Zeroizing" the switch's key automatically disables SSH (sets **ip ssh** to no). Thus, if you zeroize the key and then generate a new key, you must also re-enable SSH with the **ip ssh** command before the switch can resume SSH operation.

Configuring Key Lengths

The **crypto key generate ssh** command allows you to specify the type and length of the generated host key. The size of the host key is platform-dependent as different switches have different amounts of processing power. The size is represented by the *<keysize>* parameter and has the values shown in Table 6-2. The default value is used if **keysize** is not specified.

Table 6-2. RSA/DSA Values

Platform	Maximum RSA Key Size (in bits)	DSA Key Size (in bits)
2610	3072 Default: 1024	1024

3. Providing the Switch's Public Key to Clients

When an SSH client contacts the switch for the first time, the client will challenge the connection unless you have already copied the key into the client's "known host" file. Copying the switch's key in this way reduces the chance that an unauthorized device can pose as the switch to learn your access passwords. The most secure way to acquire the switch's public key for distribution to clients is to use a direct, serial connection between the switch and a management device (laptop, PC, or UNIX workstation), as described below.

The public key generated by the switch consists of three parts, separated by one blank space each:




Figure 6-6. Example of a Public Key Generated by the Switch

(The generated public key on the switch is always 896 bits.)

With a direct serial connection from a management station to the switch:

1. Use a terminal application such as HyperTerminal to display the switch's public key with the **show crypto host-public-key** command (figure 6-5).
2. Bring up the SSH client's "known host" file in a text editor such as Notepad as straight ASCII text, and copy the switch's public key into the file.
3. Ensure that there are no changes or breaks in the text string. (A public key must be an unbroken ASCII string. Line breaks are not allowed. Changes in the line breaks will corrupt the Key.) For example, if you are using Windows® Notepad, ensure that **Word Wrap** (in the **Edit** menu) is disabled, and that the key text appears on a single line.




Figure 6-7. Example of a Correctly Formatted Public Key

4. Add any data required by your SSH client application. For example Before saving the key to an SSH client's "known hosts" file you may have to insert the switch's IP address:

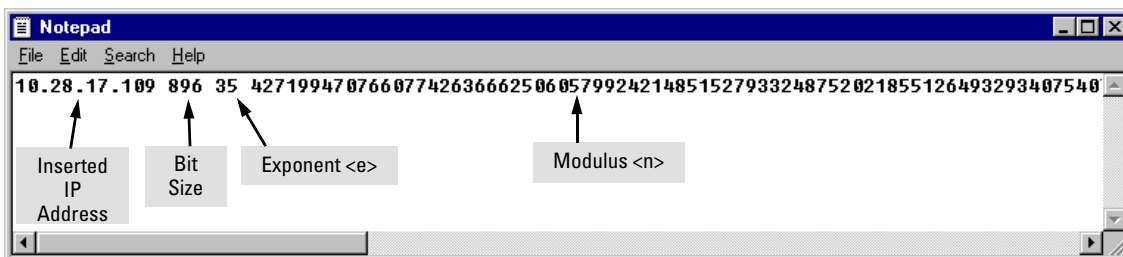


Figure 6-8. Example of a Switch Public Key Edited To Include the Switch's IP Address

For more on this topic, refer to the documentation provided with your SSH client application.

Displaying the Public Key. The switch provides three options for displaying its public key. This is helpful if you need to visually verify that the public key the switch is using for authenticating itself to a client matches the copy of this key in the client's "known hosts" file:

- **Non-encoded ASCII numeric string:** Requires a client ability to display the keys in the “known hosts” file in the ASCII format. This method is tedious and error-prone due to the length of the keys. (See figure 6-7 on page 6-14.)
- **Phonetic hash:** Outputs the key as a relatively short series of alphabetic character groups. Requires a client ability to convert the key to this format.
- **Hexadecimal hash:** Outputs the key as a relatively short series of hexadecimal numbers. Requires a parallel client ability.

For example, on the switch, you would generate the phonetic and hexadecimal versions of the switch’s public key in figure 6-7 as follows:

```
ProCurve(config)# show crypto host-public-key babble
896 xozik-kobaf-daroh-fygaz-byveb-bymiz-nupap-povaz-cesin-kafec-rixux
    host_ssh1
896 xefes-hikot-kyher-cukuz-balah-gezos-gumym-rezif-horib-cicyp-poxyx
    host_ssh2.pub
ProCurve(config)# show crypto host-public-key fingerprint
896 53:c0:14:75:72:84:90:cc:c8:ba:5e:ca:92:fc:c7:5c host_ssh1
896 bf:fb:8a:d0:10:5a:48:57:61:f9:8a:6a:61:13:8a:fb host_ssh2.pub
```

Phonetic "Hash" of Switch's Public Key

Hexadecimal "Fingerprints" of the Same Switch

Figure 6-9. Examples of Visual Phonetic and Hexadecimal Conversions of the Switch’s Public Key

The two commands shown in figure 6-9 convert the displayed format of the switch’s (host) public key for easier visual comparison of the switch’s public key to a copy of the key in a client’s “known host” file. The switch has only one RSA host key. The 'babble' and 'fingerprint' options produce two hashes for the key—one that corresponds to the challenge hash you will see if connecting with a v1 client, and the other corresponding to the hash you will see if connecting with a v2 client. These hashes do not correspond to different keys, but differ only because of the way v1 and v2 clients compute the hash of the same RSA key. The switch always uses ASCII version (without babble or fingerprint conversion) of its public key for file storage and default display format.

4. Enabling SSH on the Switch and Anticipating SSH Client Contact Behavior

The **ip ssh** command enables or disables SSH on the switch and modifies parameters the switch uses for transactions with clients. After you enable SSH, the switch can authenticate itself to SSH clients.

Note

Before enabling SSH on the switch you must generate the switch's public/private key pair. If you have not already done so, refer to "2. Generating the Switch's Public and Private Key Pair" on page 6-10.

When configured for SSH, the switch uses its host public-key to authenticate itself to SSH clients. If you also want SSH clients to authenticate themselves to the switch you must configure SSH on the switch for client public-key authentication at the login (Operator) level. To enhance security, you should also configure local, TACACS+, or RADIUS authentication at the enable (Manager) level.

Refer to "5. Configuring the Switch for SSH Authentication" on page 6-20.

SSH Client Contact Behavior. At the first contact between the switch and an SSH client, if the switch's public key has not been copied into the client, then the client's first connection to the switch will question the connection and, for security reasons, provide the option of accepting or refusing. If it is safe to assume that an unauthorized device is not using the switch's IP address in an attempt to gain access to the client's data or network, the connection can be accepted. (As a more secure alternative, the client can be directly connected to the switch's serial port to download the switch's public key into the client. See the following Note.)

Note

When an SSH client connects to the switch for the first time, it is possible for a "man-in-the-middle" attack; that is, for an unauthorized device to pose undetected as the switch, and learn the usernames and passwords controlling access to the switch. This possibility can be removed by directly connecting the management station to the switch's serial port, using a **show** command to display the switch's public key, and copying the key from the display into a file. This requires a knowledge of where the client stores public keys, plus the knowledge of what key editing and file format might be required by the client application. However, if the first contact attempt between a client and the switch does not pose a security problem, this is unnecessary.

To enable SSH on the switch.

1. Generate a public/private key pair if you have not already done so. (Refer to "2. Generating the Switch's Public and Private Key Pair" on page 6-10.)
2. Execute the **ip ssh** command.

To disable SSH on the switch, do either of the following:

- Execute **no ip ssh**.
- Zeroize the switch's existing key pair. (page 6-11).

Syntax: [no] ip ssh

Enables or disables SSH on the switch.

[cipher <cipher-type>]

Specify a cipher type to use for connection.

Valid types are:

- *aes128-cbc*
- *3des-cbc*
- *aes192-cbc*
- *aes256-cbc*
- *rijndael-cbc@lysator.liu.se*
- *aes128-ctr*
- *aes192-ctr*
- *aes256-ctr*

Default: All cipher types are available.

*Use the **no** form of the command to disable a cipher type.*

[filetransfer]

Enable/disable secure file transfer capability. SCP and SFTP secure file transfer will not function unless SSH is also enabled.

[ip-version <4|6|4or6>]

Select the IP mode to run in. The mode "ip-version 4" only accepts connections from IPv4 clients. The mode "ip-version 6" only accepts connections from IPv6 clients. The mode "ip-version 4or6" accepts connections from both IPv4 and IPv6 clients.

Default: ip-version 4or6

[mac <mac-type>]

*Allows configuration of the set of MACs that can be selected.
Valid types are:*

- *hmac-md5*
- *hmac-sha1*
- *hmac-sha1-96*
- *hmac-md5-96*

Default: All MAC types are available.

*Use the **no** form of the command to disable a MAC type.*

[port < 1-65535 | default >]

The TCP port number for SSH connections (default: 22).

Important: See “Note on Port Number” on page 6-19.

[timeout < 5 - 120 >]

*Sets the maximum length of time (in seconds) allowed
for initial protocol negotiation and authentication.*

Default: 120 seconds

[listen <oobm|data|both>]

*The **listen** parameter is available only on switches that have
a separate out-of-band management port. Values for this
parameter are:*

- **oobm** — *inbound SSH access is enabled only on
the out-of-band management port.*
- **data** — *inbound SSH access is enabled only on
the data ports.*
- **both** — *inbound SSH access is enabled on both
the out-of-band management port and on the
data ports. This is the default value.*

*Refer to Appendix G, “Network Out-of-Band Management” in
the Management and Configuration Guide for more informa-
tion on out-of-band management.*

*The **listen** parameter is not available on switches that do not
have a separate out-of-band management port.*

Note on Port Number

ProCurve recommends using the default TCP port number (22). However, you can use **ip ssh port** to specify any TCP port for SSH connections except those reserved for other purposes. Examples of reserved IP ports are 23 (Telnet) and 80 (http). Some other reserved TCP ports on the switch are 49, 80, 1506, and 1513.

ProCurve(config) ip ssh

Enable SSH

ProCurve(config)# show ip ssh

SSH Enabled : Yes

Secure Copy Enabled : No

TCP Port Number : 22

Timeout (sec) : 120

IP Version : IPv4orIPv6

Host Key Type : RSA

Host Key Size : 1024

Ciphers : aes128-cbc,3des-cbc,aes192-cbc,aes256-cbc,rijndael-cbc@lysator.liu.se,aes128-ctr,aes192-ctr,aes256-ctr

MACs : hmac-md5,hmac-sha1,hmac-sha1-96,hmac-md5-96

Ses	Type	Source IP	Port
1	console		
2	telnet		
3	ssh	12.255.255.255	
4	inactive		
5	inactive		

With SSH running, the switch allows one console session and up to five other sessions (SSH and/or Telnet). Web browser sessions are also allowed, but do not appear in the **show ip ssh** listing.

Figure 6-10. Example of Enabling IP SSH and Displaying the SSH Configuration

Caution

Protect your private key file from access by anyone other than yourself. If someone can access your private key file, they can then penetrate SSH security on the switch by appearing to be you.

SSH does not protect the switch from unauthorized access via the web interface, Telnet, SNMP, or the serial port. While web and Telnet access can be restricted by the use of passwords local to the switch, if you are unsure of the security this provides, you may want to disable web-based and/or Telnet access (**no web-management** and **no telnet**). If you need to increase SNMP security, you should use SNMP version 3 only. If you need to increase the security of your web interface see the section on SSL. Another security measure is to use the Authorized IP Managers feature described in the switch's *Management and Configuration Guide*. To protect against unauthorized

access to the serial port (and the Clear button, which removes local password protection), keep physical access to the switch restricted to authorized personnel.

5. Configuring the Switch for SSH Authentication

Note that all methods in this section result in authentication of the switch's public key by an SSH client. However, only Option B, below results in the switch also authenticating the client's public key. Also, for a more detailed discussion of the topics in this section, refer to "Further Information on SSH Client Public-Key Authentication" on page 6-24

Note

ProCurve recommends that you always assign a Manager-Level (enable) password to the switch. Without this level of protection, any user with Telnet, web, or serial port access to the switch can change the switch's configuration. *Also, if you configure only an Operator password, entering the Operator password through telnet, web, ssh or serial port access enables full manager privileges.* See "1. Assigning a Local Login (Operator) and Enable (Manager) Password" on page 6-10.

Option A: Configuring SSH Access for Password-Only SSH

Authentication. When configured with this option, the switch uses its public key to authenticate itself to a client, but uses only passwords for client authentication.

Syntax: `aaa authentication ssh login < local | tacacs | radius >[< local | none >]`

*Configures a password method for the primary and secondary login (Operator) access. If you do not specify an optional secondary method, it defaults to **none**. If the primary method is **local**, the secondary method must be **none**.*

`aaa authentication ssh enable < local | tacacs | radius>[< local | none >]`

*Configures a password method for the primary and secondary enable (Manager) access. If you do not specify an optional secondary method, it defaults to **none**. If the primary method is **local**, the secondary method must be **none**.*

Option B: Configuring the Switch for Client Public-Key SSH

Authentication. If configured with this option, the switch uses its public key to authenticate itself to a client, but the client must also provide a client public-key for the switch to authenticate. This option requires the additional step of copying a client public-key file from a TFTP server into the switch. This means that before you can use this option, you must:

1. Create a key pair on an SSH client.
2. Copy the client's public key into a public-key file (which can contain up to ten client public-keys).
3. Copy the public-key file into a TFTP server accessible to the switch and download the file to the switch.

(For more on these topics, refer to “Further Information on SSH Client Public-Key Authentication” on page 6-24.)

With steps 1 - 3, above, completed and SSH properly configured on the switch, if an SSH client contacts the switch, login authentication automatically occurs first, using the switch and client public-keys. After the client gains login access, the switch controls client access to the manager level by requiring the passwords configured earlier by the **aaa authentication ssh enable** command.

Syntax: `copy tftp pub-key-file < ip-address > < filename >`

Copies a public key file into the switch.

`aaa authentication ssh login public-key`

*Configures the switch to authenticate a client public-key at the login level with an optional secondary password method (default: **none**).*

Syntax: `aaa authentication ssh enable < local | tacacs | radius > < local | none >`

*Configures a password method for the primary and secondary enable (Manager) access. If you do not specify an optional secondary method, it defaults to **none**.*

*If the primary access method is **local**, you can only specify **none** for a secondary access method.*

Note: *The configuration of SSH clients' public keys is stored in flash memory on the switch. You also can save SSH client public-key configurations to a configuration file by entering the following commands:*

include-credentials
write memory

For more information about saving security credentials, see "Saving Security Credentials in a Config File" on page 2-10 in this guide.

For example, assume that you have a client public-key file named Client-Keys.pub (on a TFTP server at 10.33.18.117) ready for downloading to the switch. For SSH access to the switch you want to allow only clients having a private key that matches a public key found in Client-Keys.pub. For Manager-level (enable) access for successful SSH clients you want to use TACACS+ for primary password authentication and **local** for secondary password authentication, with a Manager username of "leader" and a password of "m0ns00n". To set up this operation you would configure the switch in a manner similar to the following:

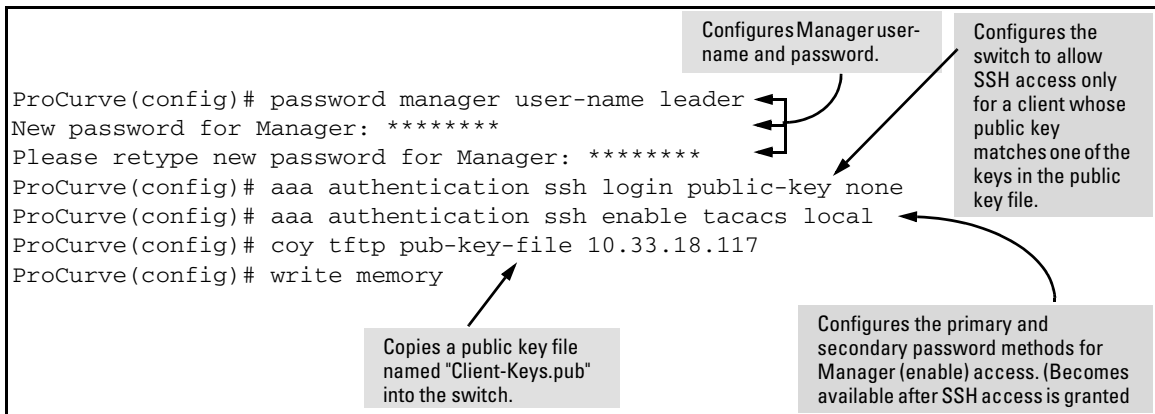


Figure 6-11. Configuring for SSH Access Requiring a Client Public-Key Match and Manager Passwords

Figure 6-12 shows how to check the results of the above commands.

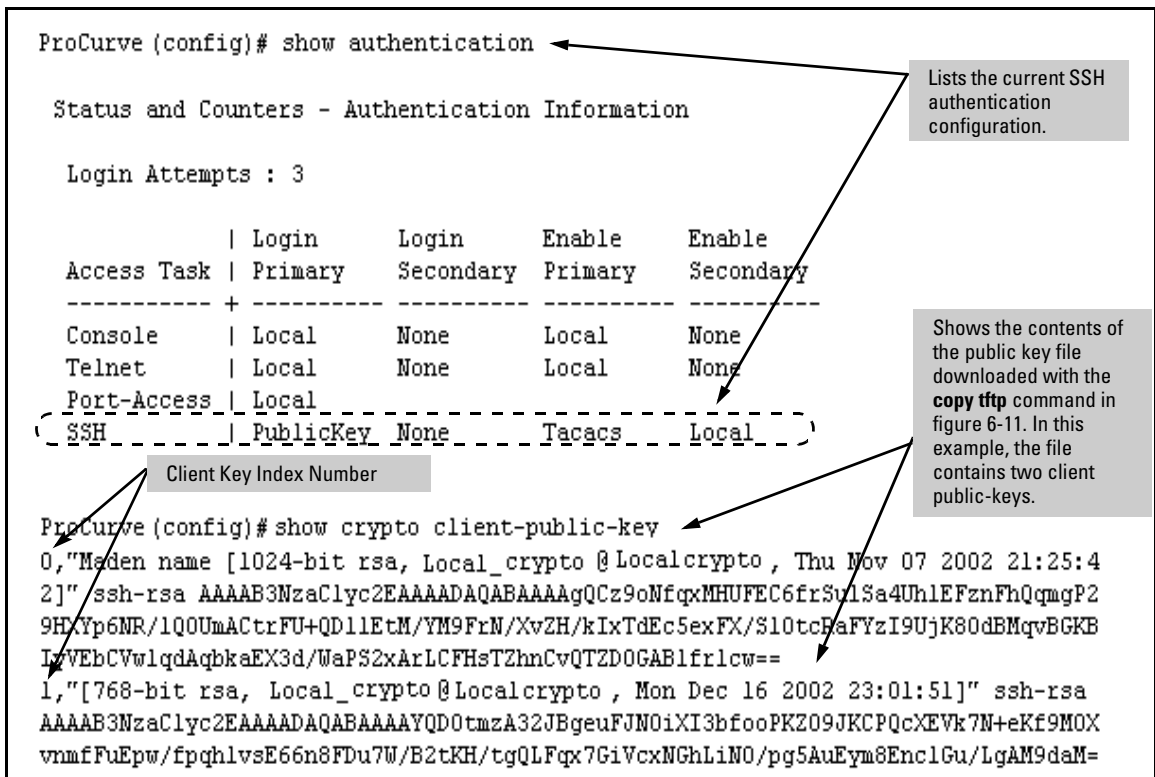


Figure 6-12. SSH Configuration and Client-Public-Key Listing From Figure 6-11

6. Use an SSH Client To Access the Switch

Test the SSH configuration on the switch to ensure that you have achieved the level of SSH operation you want for the switch. If you have problems, refer to "RADIUS-Related Problems" in the Troubleshooting chapter of the *Management and Configuration Guide* for your switch.

Further Information on SSH Client Public-Key Authentication

The section titled "5. Configuring the Switch for SSH Authentication" on page 6-20 lists the steps for configuring SSH authentication on the switch. However, if you are new to SSH or need more details on client public-key authentication, this section may be helpful.

When configured for SSH operation, the switch automatically attempts to use its own host public-key to authenticate itself to SSH clients. To provide the optional, opposite service—client public-key authentication to the switch—you can configure the switch to store up to ten public keys for authenticating clients. This requires storing an ASCII version of each client's public key (without babble conversion, or fingerprint conversion) in a client public-key file that you create and TFTP-copy to the switch. In this case, only clients that have a private key corresponding to one of the stored public keys can gain access to the switch using SSH. *That is, if you use this feature, only the clients whose public keys are in the client public-key file you store on the switch will have SSH access to the switch over the network.* If you do not allow secondary SSH login (Operator) access via local password, then the switch will refuse other SSH clients.

SSH clients that support client public-key authentication normally provide a utility to generate a key pair. The private key is usually stored in a password-protected file on the local host; the public key is stored in another file and is not protected.

(Note that even without using client public-key authentication, you can still require authentication from whoever attempts to access the switch from an SSH client— by employing the local username/password, TACACS+, or RADIUS features. Refer to "5. Configuring the Switch for SSH Authentication" on page 6-20.)

If you enable client public-key authentication, the following events occur when a client tries to access the switch using SSH:

1. The client sends its public key to the switch with a request for authentication.
2. The switch compares the client's public key to those stored in the switch's client-public-key file. (As a prerequisite, you must use the switch's **copy tftp** command to download this file to flash.)
3. If there is not a match, and you have not configured the switch to accept a login password as a secondary authentication method, the switch denies SSH access to the client.
4. If there is a match, the switch:
 - a. Generates a random sequence of bytes.
 - b. Uses the client's public key to encrypt this sequence.
 - c. Send these encrypted bytes to the client.
5. The client uses its private key to decrypt the byte sequence.
6. The client then:
 - a. Combines the decrypted byte sequence with specific session data.
 - b. Uses a secure hash algorithm to create a hash version of this information.
 - c. Returns the hash version to the switch.
7. The switch computes its own hash version of the data from step 6 and compares it to the client's hash version. If they match, then the client is authenticated. Otherwise, the client is denied access.

Using client public-key authentication requires these steps:

1. Generate a public/private key pair for each client you want to have SSH access to the switch. This can be a separate key for each client or the same key copied to several clients.
2. Copy the public key for each client into a client-public-key text file.
3. Use **copy tftp** to copy the client-public-key file into the switch. Note that the switch can hold 10 keys. The new key is appended to the client public-key file
4. Use the **aaa authentication ssh** command to enable client public-key authentication.

To Create a Client-Public-Key Text File. These steps describe how to copy client-public-keys into the switch for challenge-response authentication, and require an understanding of how to use your SSH client application.




Figure 6-13. Example of a Client Public Key

Notes

Comments in public key files, such as **smith@support.cairns.com** in figure 6-13, may appear in a SSH client application’s generated public key. While such comments may help to distinguish one key from another, they do not pose any restriction on the use of a key by multiple clients and/or users.

Public key illustrations such as the key shown in figure 6-13 usually include line breaks as a method for showing the whole key. However, in practice, line breaks in a public key will cause errors resulting in authentication failure.

1. Use your SSH client application to create a public/private key pair. Refer to the documentation provided with your SSH client application for details. The switch supports the following client-public-key properties:

Property	Supported Value	Comments
Key Format	ASCII	See figure 6-7 on page 6-14. The key must be one unbroken ASCII string. If you add more than one client-public-key to a file, terminate each key (except the last one) with a <CR><LF>. Spaces are allowed within the key to delimit the key’s components. Note that, unlike the use of the switch’s public key in an SSH client application, the format of a client-public-key used by the switch does not include the client’s IP address.
Key Type	RSA or DSA	You can choose either RSA or DSA key types when using the crypto key generate ssh command. The cert and autorun parameters only use RSA key types.
Maximum Supported Public Key Length	3072 bits	Shorter key lengths allow faster operation, but also mean diminished security.
Maximum Host Key Sizes In Bits	RSA: 1024, 2048, 3072 DSA: 1024	Includes the bit size, public index, modulus, any comments, <CR>, <LF>, and all blank spaces. If necessary, you can use an editor application to verify the size of a key. For example, placing a client-public-key into a Word for Windows text file and clicking on File Properties Statistics , lets you view the number of characters in the file, including spaces.

2. Copy the client's public key into a text file (*filename.txt*). (For example, you can use the Notepad editor included with the Microsoft® Windows® software. If you want several clients to use client public-key authentication, copy a public key for each of these clients (up to ten) into the file. Each key should be separated from the preceding key by a <CR><LF>.
3. Copy the client-public-key file into a TFTP server accessible to the switch.

Copying a client-public-key into the switch requires the following:

- One or more client-generated public keys. Refer to the documentation provided with your SSH client application.
- A copy of each client public key (up to ten) stored in a single text file or individually on a TFTP server to which the switch has access. Terminate all client public-keys in the file except the last one with a <CR><LF>.

Note on Public Keys

The actual content of a public key entry in a public key file is determined by the SSH client application generating the key. (Although you can manually add or edit any comments the client application adds to the end of the key, such as the **smith@support.cairns.com** at the end of the key in figure 6-13 on page 6-26.)

Syntax: `copy tftp pub-key-file <ip-address> <filename> [<append | manager | operator>] [oobm]`

Copies a public key file from a TFTP server into flash memory in the switch.

*The **append** option adds the key(s) for operator access.*

*The **manager** option replaces the key(s) for manager access; follow with the 'append' option to add the key(s).*

The **operator** option replaces the key(s) for operator access (default); follow with the 'append' option to add the key(s).

For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the traffic will go through the out-of-band management interface. If this parameter is not specified, the traffic goes through the data interface. The **oobm** parameter is not available on switches that do not have a separate out-of-band management port. Refer to Appendix G, "Network Out-of-Band Management" in the Management and Configuration Guide for more information on out-of-band management.

show crypto client-public-key [<manager | operator>] [keylist-str] [babble | fingerprint]

Displays the client public key(s) in the switch's current client-public-key file.

The **babble** option converts the key data to phonetic hashes that are easier for visual comparisons.

The **fingerprint** option converts the key data to hexadecimal hashes that are for the same purpose.

The **keylist-str** selects keys to display (comma-delimited list).


The **manager** option allows you to select manager public keys

The **operator** option allows you to select operator public keys.

Note

copy usb pub-key file can also be used as a method for copying a public key file to the switch.

For example, if you wanted to copy a client public-key file named **clientkeys.txt** from a TFTP server at 10.38.252.195 and then display the file contents:



```
ProCurve(config)# copy tftp pub-key-file 10.38.252.195 Clientkeys.txt
ProCurve(config)# show crypto client-public-key
0."Maden name [1024-bit rsa, Jamie_wilson@Jamiewilson, Thu Nov 07 2002 21:25:4
2]" ssh-rsa AAAAB3NzaClyc2EAAAADAQABAAQgQCz9oNfqxMHUFEC6frSulSa4Uh1EFznFhQqmgP2
9HXVp6NR/1QOUmACtrFU+QD11EtM/YM9FrN/XvZH/kIXtdEc5exFX/S10tcRaFYzI9UjK80dBMqvBGKB
IyVEbCVwlqdAqbkaEX3d/WaPS2xArLCFHsTZhnCvQTZD0GABlfrlcw==
1."[768-bit rsa, Jamie_wilson@Jamiewilson, Mon Dec 16 2002 23:01:51]" ssh-rsa
AAAAB3NzaClyc2EAAAADAQABAAQgYQD0tmzA32JBgeuFJN0iXI3bfooPKZ09JKCPQcXEVk7N+eKf9M0X
vnmfFuBpw/fpqhlvsE66n8FDu7W/B2tKH/tqQLFqx7GiVcxNGhLiN0/pq5AuEym8EnclGu/LgAM9daM=
```

Key Index Number

Figure 6-14. Example of Copying and Displaying a Client Public-Key File Containing Two Different Client Public Keys for the Same Client

Replacing or Clearing the Public Key File. The client public-key file remains in the switch's flash memory even if you erase the startup-config file, reset the switch, or reboot the switch.

- You can remove the existing client public-key file or specific keys by executing the **clear crypto public-key** command. This clears the public keys from both management modules. The module that is not active must be in standby mode.

Syntax: clear crypto public-key
Deletes the client-public-key file from the switch.

Syntax: clear crypto public-key 3
Deletes the entry with an index of 3 from the client-public-key file on the switch.

Enabling Client Public-Key Authentication. After you TFTP a client-public-key file into the switch (described above), you can configure the switch to allow the following:

- If an SSH client's public key matches the switch's client-public-key file, allow that client access to the switch. If there is not a public-key match, then deny access to that client.

Syntax: aaa authentication ssh login public-key none

Allows SSH client access only if the switch detects a match between the client's public key and an entry in the client-public-key file most recently copied into the switch.

Caution

To enable client public-key authentication to block SSH clients whose public keys are not in the client-public-key file copied into the switch, you must configure the Login Secondary as **none**. Otherwise, the switch allows such clients to attempt access using the switch's Operator password.

Messages Related to SSH Operation

Message	Meaning
00000K Peer unreachable.	File transfer did not occur. Indicates an error in communicating with the tftp server or not finding the file to download. Causes include such factors as: <ul style="list-style-type: none">• Incorrect IP configuration on the switch• Incorrect IP address in the command• Case (upper/lower) error in the filename used in the command• Incorrect configuration on the TFTP server• The file is not in the expected location.• Network misconfiguration• No cable connection to the network
00000K Transport error.	File transfer did not occur. Indicates the switch experienced a problem when trying to copy tftp the requested file. The file may not be in the expected directory, the filename may be misspelled in the command, or the file permissions may be wrong.
Cannot bind reserved TCP port <port-number>.	The ip ssh port command has attempted to configure a reserved TCP port. Use the default or select another port number. See "Note on Port Number" on page 6-19.

Message	Meaning
Client public key file corrupt or not found. Use 'copy tftp pub-key-file <ip-addr> <filename>' to download new file.	The client key does not exist in the switch. Use copy tftp to download the key from a TFTP server.
Download failed: overlength key in key file.	<p>The public key file you are trying to download has one of the following problems:</p> <ul style="list-style-type: none"> • A key in the file is too long. The maximum key length is 1024 characters, including spaces. This could also mean that two or more keys are merged together instead of being separated by a <CR><LF>. • There are more than ten public keys in the key file and switch total. Delete some keys from the switch or file. The switch does not detect duplicate keys. • One or more keys in the file is corrupted or is not a valid public key. <p>Refer to “To Create a Client-Public-Key Text File” on page 26 for information on client-public-key properties.</p>
Download failed: too many keys in key file.	
Download failed: one or more keys is not a valid public key.	
Error: Requested keyfile does not exist.	The client key does not exist in the switch. Use copy tftp to download the key from a TFTP server.
Generating new RSA host key. If the cache is depleted, this could take up to two minutes.	After you execute the generate ssh [dsa rsa] command, the switch displays this message while it is generating the key.
Host RSA key file corrupt or not found. Use ' generate ssh [dsa rsa]' to create new host key.	The switch's key is missing or corrupt. Use the generate ssh [dsa rsa] command to generate a new key for the switch.

Logging Messages

There are event log messages when a new key is generated and zeroized for the server:

```
ssh: New <num-bits> -bit [rsa | dsa] SSH host key installed
ssh: SSH host key zeroized
```

There are also messages that indicates when a client public key is installed or removed:

```
ssh: <num-bits>-bit [rsa | dsa] client public key [installed | removed] ([manager|
operator] access) (key_comment)
```

Note: Only up to 39 characters of the key comment are included in the event log message.

Debug Logging

To add ssh messages to the debug log output, enter this command:

```
ProCurve# debug ssh LOGLEVEL
```

where LOGLEVEL is one of the following (in order of increasing verbosity):

- fatal
- error
- info
- verbose
- debug
- debug2
- debug3

Configuring Secure Socket Layer (SSL)

Contents

Overview	7-2
Terminology	7-3
Prerequisite for Using SSL	7-5
Steps for Configuring and Using SSL for Switch and Client Authentication	7-5
General Operating Rules and Notes	7-6
Configuring the Switch for SSL Operation	7-7
1. Assigning a Local Login (Operator) and Enabling (Manager) Password	7-7
2. Generating the Switch's Server Host Certificate	7-8
To Generate or Erase the Switch's Server Certificate with the CLI	7-9
Comments on Certificate Fields.	7-10
Generate a Self-Signed Host Certificate with the Web browser interface	7-12
Generate a CA-Signed server host certificate with the Web browser interface	7-15
3. Enabling SSL on the Switch and Anticipating SSL Browser Contact Behavior	7-17
Using the CLI Interface to Enable SSL	7-19
Using the Web Browser Interface to Enable SSL	7-19
Common Errors in SSL setup	7-21

Overview

Feature	Default	Menu	CLI	Web
Generating a Self Signed Certificate on the switch	No	n/a	page 7-8	page 7-12
Generating a Certificate Request on the switch	No	n/a	n/a	page 7-15
Enabling SSL	Disabled	n/a	page 7-17	page 7-19

The switches covered in this guide use Secure Socket Layer Version 3 (SSLv3) and support for Transport Layer Security(TLSv1) to provide remote web access to the switches via encrypted paths between the switch and management station clients capable of SSL/TLS operation.

Note

ProCurve Switches use SSL and TLS for all secure web transactions, and all references to SSL mean using one of these algorithms unless otherwise noted

SSL provides all the web functions but, unlike standard web access, SSL provides encrypted, authenticated transactions. The authentication type includes server certificate authentication with user password authentication.

Note

SSL in the switches covered in this guide is based on the OpenSSL software toolkit. For more information on OpenSSL, visit www.openssl.com.

Server Certificate authentication with User Password

Authentication . This option is a subset of full certificate authentication of the user and host. It occurs only if the switch has SSL enabled. As in figure 7-1, the switch authenticates itself to SSL enabled web browser. Users on SSL browser then authenticate themselves to the switch (operator and/or manger levels) by providing passwords stored locally on the switch or on a TACACS+ or RADIUS server. However, the client does not use a certificate to authenticate itself to the switch.

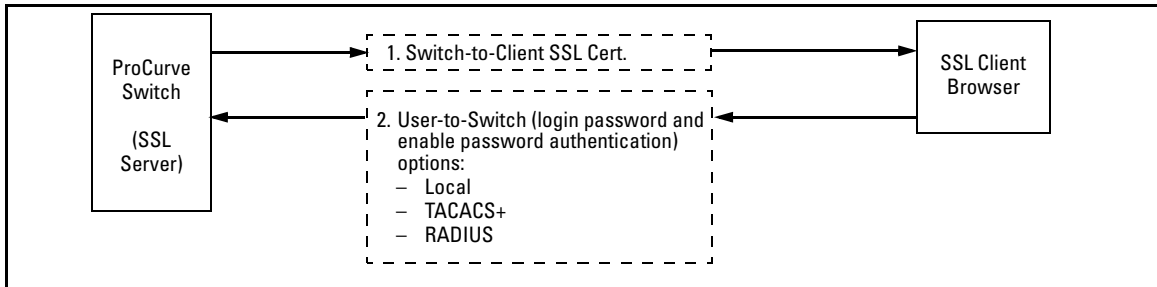


Figure 7-1. Switch/User Authentication

SSL on the switches covered in this guide supports these data encryption methods:

- 3DES (168-bit, 112 Effective)
- DES (56-bit)
- RC4 (40-bit, 128-bit)

Note:

ProCurve Switches use RSA public key algorithms and Diffie-Hellman, and all references to a key mean keys generated using these algorithms unless otherwise noted

Terminology

- **SSL Server:** An ProCurve switch with SSL enabled.
- **Key Pair:** Public/private pair of RSA keys generated by switch, of which public portion makes up part of server host certificate and private portion is stored in switch flash (not user accessible).
- **Digital Certificate:** A certificate is an electronic “passport” that is used to establish the credentials of the subject to which the certificate was issued. Information contained within the certificate includes: name of the subject, serial number, date of validity, subject's public key, and the digital signature of the authority who issued the certificate. Certificates on ProCurve switches conform to the X.509v3 standard, which defines the format of the certificate.
- **Self-Signed Certificate:** A certificate not verified by a third-party certificate authority (CA). Self-signed certificates provide a reduced level of security compared to a CA-signed certificate.
- **CA-Signed Certificate:** A certificate verified by a third party certificate authority (CA). Authenticity of CA-Signed certificates can be verified by an audit trail leading to a trusted root certificate.

- **Root Certificate:** A trusted certificate used by certificate authorities to sign certificates (CA-Signed Certificates) and used later on to verify that authenticity of those signed certificates. Trusted certificates are distributed as an integral part of most popular web clients. (see browser documentation for which root certificates are pre-installed).
- **Manager Level:** Manager privileges on the switch.
- **Operator Level:** Operator privileges on the switch.
- **Local password or username:** A Manager-level or Operator-level password configured in the switch.
- **SSL Enabled:** (1)A certificate key pair has been generated on the switch (web interface or CLI command: **crypto key generate cert [key size]**) (2) A certificate been generated on the switch (web interface or CLI command: **crypto host-cert generate self-signed [arg-list]**) and (3) SSL is enabled (web interface or CLI command: **web-management ssl**). (You can generate a certificate without enabling SSL, but you cannot enable SSL without first generating a Certificate.

Prerequisite for Using SSL

Before using the switch as an SSL server, you must install a publicly or commercially available SSL enabled web browser application on the computer(s) you use for management access to the switch.

Steps for Configuring and Using SSL for Switch and Client Authentication

The general steps for configuring SSL include:

A. Client Preparation

1. Install an SSL capable browser application on a management station you want to use for access to the switch. (Refer to the documentation provided with your browser.)

Note:

The latest versions of Microsoft Internet Explorer and Netscape web browser support SSL and TLS functionality. See browser documentation for additional details

B. Switch Preparation

1. Assign a login (Operator) and enable (Manager) password on the switch. (page 7-7)
2. Generate a host certificate on the switch. (page 7-8)
 - i. Generate certificate key pair
 - ii. Generate host certificate

You need to do this only once. The switch's own public/private certificate key pair and certificate are stored in the switch's flash memory and are not affected by reboots or the erase startup-config command. (You can remove or replace this certificate, if necessary.) The certificate key pair and the SSH key pair are independent of each other, which means a switch can have two keys pairs stored in flash.

3. Enable SSL on the switch. (page 7-17)
4. Use your SSL enabled browser to access the switch using the switch's IP address or DNS name (if allowed by your browser). Refer to the documentation provided with the browser application.

General Operating Rules and Notes

- Once you generate a certificate on the switch you should avoid re-generating the certificate without a compelling reason. Otherwise, you will have to re-introduce the switch's certificate on all management stations (clients) you previously set up for SSL access to the switch. In some situations this can temporarily allow security breaches.
- The switch's own public/private certificate key pair and certificate are stored in the switch's flash memory and are not affected by reboots or the erase startup-config command
- The public/private certificate key pair is not be confused with the SSH public/private key pair. The certificate key pair and the SSH key pair are independent of each other, which means a switch can have two keys pairs stored in flash.

Configuring the Switch for SSL Operation

SSL-Related CLI Commands in This Section	Page
web-management ssl	page 7-19
show config	page 7-19
show crypto host-cert	page 7-12
crypto key	
generate cert [rsa] <512 768 1024>	page 7-10
zeroize cert	page 7-10
crypto host-cert	
generate self-signed [arg-list]	page 7-10
zeroize	page 7-10

1. Assigning a Local Login (Operator) and Enabling (Manager) Password

At a minimum, ProCurve recommends that you always assign at least a Manager password to the switch. Otherwise, under some circumstances, anyone with Telnet, web, or serial port access could modify the switch's configuration.

Using the web browser interface To Configure Local Passwords. You can configure both the Operator and Manager password on one screen. To access the web browser interface, refer to the chapter titled "Using the ProCurve Web Browser Interface" in the *Management and Configuration Guide* for your switch.

Configuring Secure Socket Layer (SSL)

Configuring the Switch for SSL Operation

The screenshot shows the HP ProCurve Switch web interface. At the top, the title bar reads "HP ProCurve Switch - Status: Information" and "HP JXXXX ProCurve Switch". Below the title bar is a navigation menu with tabs: Identity, Status, Configuration, Security (highlighted), Diagnostics, and Support. Under the Security tab, there are sub-tabs: Device Passwords (highlighted), Authorized Addresses, Port Security, Intrusion Log, and SSL. The main content area is titled "Security Tab" and contains two sections: "Read-Only Access" and "Read-Write Access". The "Read-Only Access" section has three input fields: "Operator User Name:", "Operator Password:", and "Confirm Operator Password:". The "Read-Write Access" section has three input fields: "Manager User Name:", "Manager Password:", and "Confirm Manager Password:". At the bottom of the page are two buttons: "Apply Changes" and "Clear Changes". A callout box labeled "Password Button" points to the "Device Passwords" sub-tab. Another callout box labeled "Security Tab" points to the "Security" main tab.

Figure 7-2. Example of Configuring Local Passwords

1. Proceed to the security tab and select device passwords button.
2. Click in the appropriate box in the Device Passwords window and enter user names and passwords. You will be required to repeat the password strings in the confirmation boxes.

Both the user names and passwords can be up to 16 printable ASCII characters.
3. Click on **[Apply Changes]** button to activate the user names and passwords.

2. Generating the Switch's Server Host Certificate

You must generate a server certificate on the switch before enabling SSL. The switch uses this server certificate, along with a dynamically generated session key pair to negotiate an encryption method and session with a browser trying

to connect via SSL to the switch. (The session key pair mentioned above is not visible on the switch. It is a temporary, internally generated pair used for a particular switch/client session, and then discarded.)

The server certificate is stored in the switch's flash memory. The server certificate should be added to your certificate folder on the SSL clients who you want to have access to the switch. Most browser applications automatically add the switch's host certificate to their certificate folder on the first use. This method does allow for a security breach on the first access to the switch. (Refer to the documentation for your browser application.)

There are two types of certificates that can be used for the switch's host certificate. The first type is a self-signed certificate, which is generated and digitally signed by the switch. Since self-signed certificates are not signed by a third-party certificate authority, there is no audit trail to a root CA certificate and no fool-proof means of verifying authenticity of certificate. The second type is a certificate authority-signed certificate, which is digitally signed by a certificate authority, has an audit trail to a root CA certificate, and can be verified unequivocally

Note:

There is usually a fee associated with receiving a verified certificate and the valid dates are limited by the root certificate authority issuing the certificate.

When you generate a certificate key pair and/or certificate on the switch, the switch places the key pair and/or certificate in flash memory (and not in running config). Also, the switch maintains the certificate across reboots, including power cycles. You should consider this certificate to be "permanent"; that is, avoid re-generating the certificate without a compelling reason. Otherwise, you will have to re-introduce the switch's host certificate on all management stations you have set up for SSL access to the switch using the earlier certificate.

Removing (zeroizing) the switch's certificate key pair or certificate render the switch unable to engage in SSL operation and automatically disables SSL on the switch. (To verify whether SSL is enabled, execute **show config**.)

To Generate or Erase the Switch's Server Certificate with the CLI

Because the host certificate is stored in flash instead of the running-config file, it is not necessary to use **write memory** to save the certificate. Erasing the host certificate automatically disables SSL.

CLI commands used to generate a Server Host Certificate.

Syntax: `crypto key generate cert [rsa] < 512 | 768 | 1024 >`

Generates a key pair for use in the certificate.

`crypto key zeroize cert`

Erases the switch's certificate key and disables SSL operation.

`crypto host-cert generate self-signed [arg-list]`

Generates a self signed host certificate for the switch. If a switch certificate already exists, replaces it with a new certificate. (See the Note, above.)

`crypto host-cert zeroize`

Erases the switch's host certificate and disables SSL operation.

To generate a host certificate from the CLI:

- i. Generate a certificate key pair. This is done with the **crypto key generate cert** command. The default key size is 512.

Note:

If a certificate key pair is already present in the switch, it is not necessary to generate a new key pair when generating a new certificate. The existing key pair may be re-used and the `crypto key generate cert` command does not have to be executed

- ii. Generate a new self-signed host certificate. This is done with the **crypto host-cert generate self-signed [Arg-List]** command.

Note

When generating a self-signed host certificate on the CLI if there is not certificate key generated this command will fail.

Comments on Certificate Fields.

There are a number arguments used in the generation of a server certificate. table 7-1, "Certificate Field Descriptions" describes these arguments.

Table 7-1. Certificate Field Descriptions

Field Name	Description
Valid Start Date	This should be the date you desire to begin using the SSL functionality.
Valid End Date	This can be any future date, however good security practices would suggest a valid duration of about one year between updates of passwords and keys.
Common name	This should be the IP address or domain name associated with the switch. Your web browser may warn you if this field does not match the URL entered into the web browser when accessing the switch
Organization	This is the name of the entity (e.g. company) where the switch is in service.
Organizational Unit	This is the name of the sub-entity (e.g. department) where the switch is in service.
City or location	This is the name of the city where switch is in service
State name	This is the name of the state or province where switch is in service
Country code	This is the ISO two-letter country-code where switch is in service

For example, to generate a key and a new host certificate:

```
ProCurve(config)# crypto key generate cert 512
Installing new RSA key. If the key/entropy cache is
depleted, this could take up to a minute.
ProCurve(config)# crypto host-cert generate self-signed
Validity start date [01/01/1970]: 01/01/2002
Validity end date   [01/01/2003]: 01/01/2004
Common name        [10.255.255.255]: 10.255.255.255
Organization        [Company Name]: Hewlett Packard
Organizational unit  [Dept Name]: ProCurve Network
City or location     [City]: Roseville
State name           [State]: Ca
Country code         [US]: US
```

Generate New Key

Generate New Certificate

Enter certificate Arguments

Figure 7-3. Example of Generating a Self-Signed Server Host certificate on the CLI for the Switch.

Notes

“Zeroizing” the switch’s server host certificate or key automatically disables SSL (sets **web-management ssl** to **No**). Thus, if you zeroize the server host certificate or key and then generate a new key and server certificate, you must also re-enable SSL with the **web-management ssl** command before the switch can resume SSL operation.

CLI Command to view host certificates.

Syntax: show crypto host-cert

Displays switch's host certificate

To view the current host certificate from the CLI you use the **show crypto host-cert** command.

For example, to display the new server host certificate:

```
ProCurve(config)#show crypto host-cert ← Show host certificate command
Version: 1 (0x0)
Serial Number: 0 (0x0)
Issuer: CN=10.255.255.255, L=Roseville, ST=Ca, C=US, O=Hewlett Packard, OU=ProCurve Network
Validity
  Not Before: Jan  1 00:00:00 2002 GMT
  Not After : Jan  1 23:59:59 2004 GMT
Subject: CN=10.255.255.255, L=Roseville, ST=Ca, C=US, O=Hewlett Packard, OU=ProCurve Network
Subject Public Key Info:
  Public Key Algorithm: rsaEncryption
  RSA Public Key: (512 bit)
    Modulus (512 bit):
      00:db:18:4b:ce:3e:7d:5a:90:d8:a5:50:d5:2a:e9:
      60:78:d1:35:82:e9:27:71:5d:45:8d:0a:b9:b4:55:
      65:c7:d1:1c:4e:30:5e:20:a6:2d:62:9c:4c:cd:40:
      a0:6a:0b:cb:1c:ce:90:1c:2c:ad:26:fc:0b:07:ae:
      db:11:65:d6:47
    Exponent: 35 (0x23)
  Signature Algorithm: md5WithRSAEncryption
    d6:d0:98:6b:b9:a5:54:96:d9:be:fa:b9:99:f9:d8:6f:94:42:
    30:ea:c4:1d:88:e6:7b:19:18:22:84:f6:8c:ea:46:d7:ab:42:
    26:48:77:0c:60:57:8c:33:bc:08:d8:f7:c6:1f:ef:15:b7:24:
    f3:fa:92:b1:1f:7d:9e:c1:fd:83

MD5 Fingerprint: C969 E196 49C3 4609 AFC6 BDE1 2087 00A7
SHA1 Fingerprint: 93C7 0753 F805 26DC 4E39 EAF2 9C18 174F 7A63 E3C5
```

Figure 7-4. Example of show crypto host-cert command

Generate a Self-Signed Host Certificate with the Web browser interface

You can configure SSL from the web browser interface. For more information on how to access the web browser interface refer to the chapter titled “Using the ProCurve Web Browser Interface” in the *Management and Configuration Guide* for your switch.

To generate a self signed host certificate from the web browser interface:

- i. Proceed to the Security tab then the SSL button. The SSL configuration screen is split up into two halves. The left half is used in creating a new certificate key pair and (self-signed / CA-signed) certificate. The right half displays information on the currently installed certificate.
- ii. Select the Generate Certificate button.
- iii. Select Self signed certificate in the type box.
- iv. Select the RSA key size desired. If you do not wish to generate a new key then just select current from the list.
- v. Fill in remaining certificate arguments (refer to “To Generate or Erase the Switch’s Server Certificate with the CLI” on page 7-9).
- vi. Click on the **[Apply Changes]** button to generate a new certificate and key if selected.

Note:

When generating a self-signed host certificate, if no key is present and the current option is selected in the RSA key size box and error will be generated. New key generation can take up to two minutes if the key queue is empty.

Configuring Secure Socket Layer (SSL)

Configuring the Switch for SSL Operation

For example, to generate a new host certificate via the web browsers interface:

HP ProCurve Switch - Status: Information
HP JXXXX ProCurve Switch

Identity Status Configuration **Security** Diagnostics Support

Device Passwords Authorized Addresses Port Security Intrusion Log **SSL**

SSL Settings

SSL Enable: ☒ Off Port: 443

☒ Create Certificate/ Certificate Request ☐ Use Installed Certificate

Certificate Type: Self Signed

RSA Key Size: 512

Validity Start Date: Month Day Year

Validity End Date: Month Day Year

Common Name: 10.255.255.255

Organization Name: Company Name

Organization Unit: Department Name

City: City

State: State

Country: US - United States

Validity End Date:

Common Name :

Organization Name :

Organization Unit :

City :

State:

Country :

Fingerprint MD5:

SHA :

Apply Changes Clear Changes

Figure 7-5. Self-Signed Certificate generation via SSL Web Browser Interface Screen

To view the current host certificate in the web browser interface:

1. Proceed to the **Security** tab
2. Then the **[SSL]** button

HP ProCurve Switch - Status: Information
HP JXXXXXProCurve Switch

Identity Status Configuration **Security** Diagnostics Support

Device Passwords Authorized Addresses Port Security Intrusion Log **SSL**

SSL Settings Current SSL Host Certificate

SSL Enable: ☒ Off Port: 443

☐ Create Certificate/ Certificate Request

Certificate Type: Self Signed

RSA Key Size: 512

Certificate Information Fields

Validity Start Date: Month Day Year

Validity End Date: Month Day Year

Common Name: 10.255.255.255

Organization Name: Hewlett Packard

Organization Unit: ProCurve Network

City: Roseville

State: Ca

Country: US - United States

☒ Use Installed Certificate

Installed Certificate

Certificate Type: Self-Signed

RSA Key Size: 512 bits

Validity Start Date: 1/1/2002

Validity End Date: 1/1/2003

Common Name: 10.255.255.255

Organization Name: Hewlett Packard

Organization Unit: ProCurve Network

City: Roseville

State: Ca

Country: US

Fingerprint: BE01 E39E D49C 2575 200B 30E6 E080 38C3 CE94 BFD8 86F8 1887 BE24 F173 55D4 BE0A 4E05 2C40

SHA: BE24 F173 55D4 BE0A 4E05 2C40

Apply Changes Clear Changes

Figure 7-6. Web browser Interface showing current SSL Host Certificate

Generate a CA-Signed server host certificate with the Web browser interface

To install a CA-Signed server host certificate from the web browser interface. For more information on how to access the web browser interface, refer to the chapter titled “Using the ProCurve Web Browser Interface” in the *Management and Configuration Guide* for your switch.

The installation of a CA-signed certificate involves interaction with other entities and consists of three phases. The first phase is the creation of the CA certificate request, which is then copied off from the switch for submission to the certificate authority. The second phase is the actual submission process that involves having the certificate authority verify the certificate request and then digitally signing the request to generate a certificate response (the usable server host certificate). The third phase is the download phase consisting of pasting to the switch web server the certificate response, which is then validated by the switch and put into use by enabling SSL

To generate a certificate request from the web browser interface:

- i. Select the **Security** tab, then select the **[SSL]** button
- ii. Select the **Create Certificate/Certificate Request** radio button.
- iii. Select **Create CA Request** from the **Certificate Type** drop-down list.
- iv. Select the key size from the RSA Key Size drop-down list. If you wish to re-use the current certificate key, select **Current** from the **RSA Key Size** drop-down list.
- v. Fill in remaining certificate arguments (Refer to “Comments on Certificate Fields.” on page 7-10.)
- vi. Click on **[Apply Changes]** to create the certificate request. A new web browser page appears, consisting of two text boxes. The switch uses the upper text box for the certificate request text. The lower text box appears empty. You will use it for pasting in the certificate reply after you receive it from the certificate authority. (This authority must return a non- PEM encoded certificate request reply.
- vii. After the certificate authority processes your request and sends you a certificate reply (that is, an installable certificate), copy and paste it into the lower text box.
- viii. Click on the **[Apply Changes]** button to install the certificate.

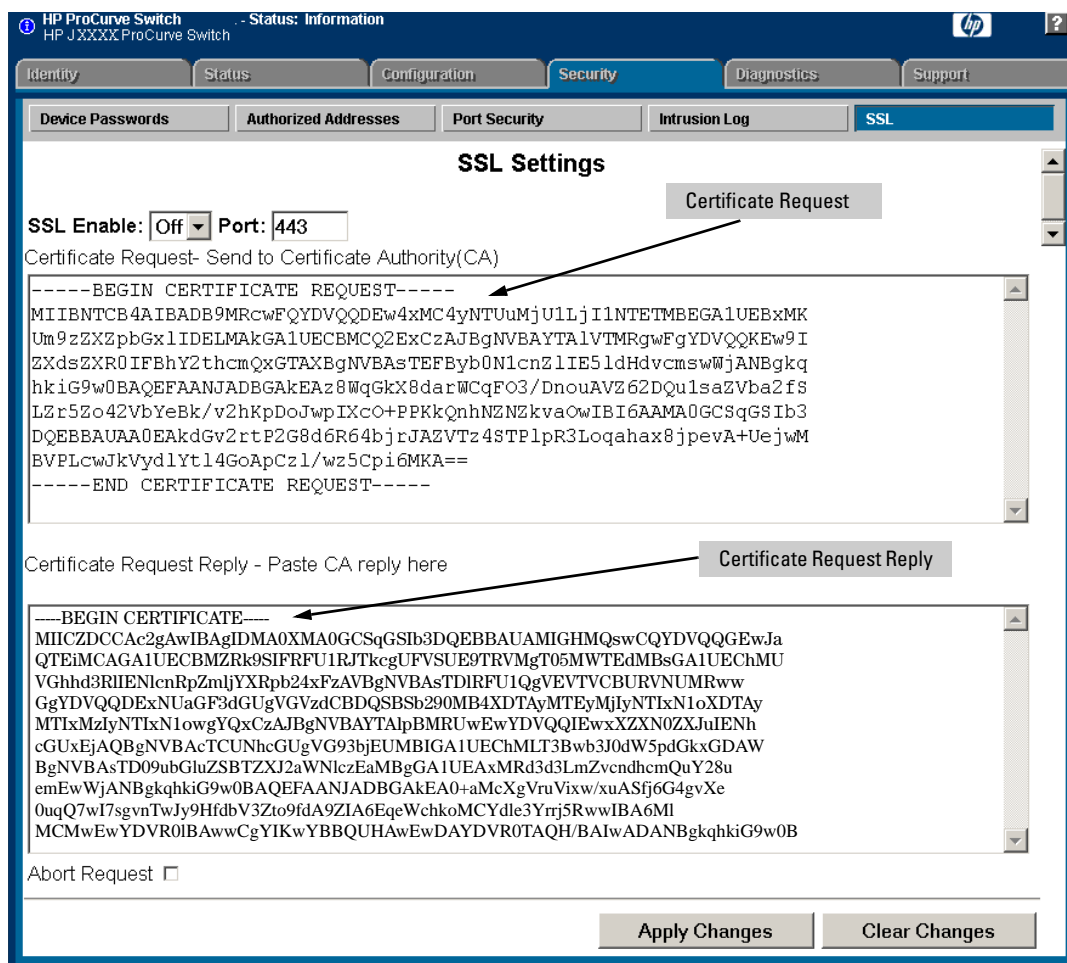


Figure 7-7. Request for Verified Host Certificate Web Browser Interface Screen

3. Enabling SSL on the Switch and Anticipating SSL Browser Contact Behavior

The **web-management ssl** command enables SSL on the switch and modifies parameters the switch uses for transactions with clients. After you enable SSL, the switch can authenticate itself to SSL enabled browsers. If you want to disable SSL on the switch, use the **no web-management ssl** command.

Note

Before enabling SSL on the switch you must generate the switch's host certificate and key. If you have not already done so, refer to "2. Generating the Switch's Server Host Certificate" on page 7-8.

When configured for SSL, the switch uses its host certificate to authenticate itself to SSL clients, however unless you disable the standard ProCurve web browser interface with the **no web-management** command it will be still available for unsecured transactions.

SSL Client Contact Behavior. At the first contact between the switch and an SSL client, if you have not copied the switch's host certificate into the browser's certificate folder, your browser's first connection to the switch will question the connection and, for security reasons, give you the option of accepting or refusing. If a CA-signed certificate is used on the switch, for which a root certificate exists on the client browser side, then the browser will NOT prompt the user to ensure the validity of the certificate. The browser will be able to verify the certificate chain of the switch server certificate up to the root certificate installed in the browser, thus authenticating the switch unequivocally. As long as you are confident that an unauthorized device is not using the switch's IP address in an attempt to gain access to your data or network, you can accept the connection.

Note

When an SSL client connects to the switch for the first time, it is possible for a "man-in-the-middle" attack; that is, for an unauthorized device to pose undetected as the switch, and learn the usernames and passwords controlling access to the switch. When using self-signed certificates with the switch, there is a possibility for a "man-in-the-middle" attack when connecting for the first time; that is, an unauthorized device could pose undetected as a switch, and learn the usernames and passwords controlling access to the switch. Use caution when connecting for the first time to a switch using self-signed certificates. Before accepting the certificate, closely verify the contents of the certificate (see browser documentation for additional information on viewing contents of certificate).

The security concern described above does not exist when using CA-signed certificates that have been generated by certificate authorities that the web browser already trusts

Using the CLI Interface to Enable SSL

Syntax: [no] web-management ssl

Enables or disables SSL on the switch.

[port < 1-65535 | default:443 >]

*The TCP port number for SSL connections (default: 443). **Important:** See “Note on Port Number” on page 7-20.*

show config

*Shows status of the SSL server. When enabled **web-management ssl** will be present in the config list.*

To enable SSL on the switch

1. Generate a Host certificate if you have not already done so. (Refer to “2. Generating the Switch’s Server Host Certificate” on page 7-8.)
2. Execute the **web-management ssl** command.

To disable SSL on the switch, do either of the following:

- Execute **no web-management ssl**.
- Zeroize the switch’s host certificate or certificate key. (page 7-9).

Using the Web Browser Interface to Enable SSL

To enable SSL on the switch

- i. Proceed to the Security tab then the SSL button
- ii. Select SSL Enable to on and enter the TCP port you desire to connect on.
- iii. Click on the **[Apply Changes]** button to enable SSL on the port.

To disable SSL on the switch, do either of the following:

- i. Proceed to the Security tab then the SSL button
- ii. Select SSL Enable to off .
- iii. Click on the **[Apply Changes]** button to enable SSL on the port.

Configuring Secure Socket Layer (SSL)

Configuring the Switch for SSL Operation




Figure 7-8. Using the web browser interface to enable SSL and select TCP port number

Note on Port Number

ProCurve recommends using the default IP port number (443). However, you can use **web-management ssl tcp-port** to specify any TCP port for SSL connections except those reserved for other purposes. Examples of reserved IP ports are 23 (Telnet) and 80 (http). Some other reserved TCP ports on the switches are 49, 80, 1506, and 1513.

Caution

SSL does not protect the switch from unauthorized access via the Telnet, SNMP, or the serial port. While Telnet access can be restricted by the use of passwords local to the switch, if you are unsure of the security this provides, you may want to disable Telnet access (**no telnet**). If you need to increase SNMP security, use SNMP version 3 only for SNMP access. Another security measure is to use the Authorized IP Managers feature described in the switch's *Security Guide*. To protect against unauthorized access to the serial port (and the Clear button, which removes local password protection), keep physical access to the switch restricted to authorized personnel.

Common Errors in SSL setup

Error During	Possible Cause
Generating host certificate on CLI	You have not generated a certificate key. (Refer to "CLI commands used to generate a Server Host Certificate" on page 7-10.)
Enabling SSL on the CLI or Web browser interface	<p>You have not generated a host certificate. (Refer to "Generate a Self-Signed Host Certificate with the Web browser interface" on page 7-12.)</p> <p>You may be using a reserved TCP port. (Refer to "Note on Port Number" on page 7-20.)</p>
Unable to Connect with SSL	<p>You may not have SSL enabled (Refer to "3. Enabling SSL on the Switch and Anticipating SSL Browser Contact Behavior" on page 7-17.)</p> <p>Your browser may not support SSLv3 or TLSv1 or it may be disabled. (Refer to the documentation provided for your browser.)</p>

Configuring Advanced Threat Protection

Contents

Introduction	8-3
DHCP Snooping	8-4
Overview	8-4
Enabling DHCP Snooping	8-5
Enabling DHCP Snooping on VLANs	8-7
Configuring DHCP Snooping Trusted Ports	8-8
Configuring Authorized Server Addresses	8-9
Using DHCP Snooping with Option 82	8-9
Changing the Remote-id from a MAC to an IP Address	8-11
Disabling the MAC Address Check	8-11
The DHCP Binding Database	8-12
Operational Notes	8-13
Log Messages	8-14
Dynamic ARP Protection	8-16
Introduction	8-16
Enabling Dynamic ARP Protection	8-18
Configuring Trusted Ports	8-18
Adding an IP-to-MAC Binding to the DHCP Database	8-20
Configuring Additional Validation Checks on ARP Packets	8-21
Verifying the Configuration of Dynamic ARP Protection	8-21
Displaying ARP Packet Statistics	8-22
Monitoring Dynamic ARP Protection	8-23
Dynamic IP Lockdown	8-23
Protection Against IP Source Address Spoofing	8-24
Prerequisite: DHCP Snooping	8-24
Filtering IP and MAC Addresses Per-Port and Per-VLAN	8-25
Enabling Dynamic IP Lockdown	8-26

Operating Notes	8-26
Adding an IP-to-MAC Binding to the DHCP Binding Database	8-28
Potential Issues with Bindings	8-28
Adding a Static Binding	8-29
Verifying the Dynamic IP Lockdown Configuration	8-29
Displaying the Static Configuration of IP-to-MAC Bindings	8-30
Debugging Dynamic IP Lockdown	8-31
Using the Instrumentation Monitor	8-33
Operating Notes	8-34
Configuring Instrumentation Monitor	8-35
Examples	8-36
Viewing the Current Instrumentation Monitor Configuration	8-37

Introduction

As your network expands to include an increasing number of mobile devices, continuous Internet access, and new classes of users (such as partners, temporary employees, and visitors), additional protection from attacks launched from both inside and outside your internal network is often necessary.

Advanced threat protection can detect port scans and hackers who try to access a port or the switch itself. The following software features provide advanced threat protection and are described in this chapter:

- DHCP snooping: Protects your network from common DHCP attacks, such as:
 - Address spoofing in which an invalid IP address or network gateway address is assigned by a rogue DHCP server.
 - Address exhaustion of available addresses in the network DHCP server, caused by repeated attacker access to the network and numerous IP address requests.
- Dynamic ARP protection: Protects your network from ARP cache poisoning as in the following cases:
 - An unauthorized device forges an illegitimate ARP response and network devices use the response to update their ARP caches.
 - A denial-of-service (DoS) attack from unsolicited ARP responses changes the network gateway IP address so that outgoing traffic is prevented from leaving the network and overwhelms network devices.
- Instrumentation monitor: Protects your network from a variety of other common attacks besides DHCP and ARP attacks, including:
 - Attempts at a port scan to expose a vulnerability in the switch, indicated by an excessive number of packets sent to closed TCP/UDP ports
 - Attempts to fill all IP address entries in the switch's forwarding table and cause legitimate traffic to be dropped, indicated by an increased number of learned IP destination addresses
 - Attempts to spread viruses, indicated by an increased number of ARP request packets

- Attempts to exhaust system resources so that sufficient resources are not available to transmit legitimate traffic, indicated by an unusually high use of specific system resources
- Attempts to attack the switch's CPU and introduce delay in system response time to new network events
- Attempts by hackers to access the switch, indicated by an excessive number of failed logins or port authentication failures
- Attempts to deny switch service by filling the forwarding table, indicated by an increased number of learned MAC addresses or a high number of MAC address moves from one port to another
- Attempts to exhaust available CPU resources, indicated by an increased number of learned MAC address events being discarded

DHCP Snooping

Command	Page
dhcp-snooping	page 8-5
authorized-server	page 8-9
database	page 8-12
option	page 8-9
trust	page 8-8
verify	page 8-11
vlan	page 8-7
show dhcp-snooping	page 8-6
show dhcp-snooping stats	page 8-6
dhcp-snooping binding	page 8-13
debug dhcp-snooping	page 8-13

Overview

You can use DHCP snooping to help avoid the Denial of Service attacks that result from unauthorized users adding a DHCP server to the network that then provides invalid configuration data to other DHCP clients on the network.

DHCP snooping accomplishes this by allowing you to distinguish between trusted ports connected to a DHCP server or switch and untrusted ports connected to end-users. DHCP packets are forwarded between trusted ports without inspection. DHCP packets received on other switch ports are inspected before being forwarded. Packets from untrusted sources are dropped. Conditions for dropping packets are shown below.

Condition for Dropping a Packet	Packet Types
A packet from a DHCP server received on an untrusted port	DHCPOFFER, DHCPACK, DHCPNACK
If the switch is configured with a list of authorized DHCP server addresses and a packet is received from a DHCP server on a trusted port with a source IP address that is not in the list of authorized DHCP server addresses.	DHCPOFFER, DHCPACK, DHCPNACK
Unless configured to not perform this check, a DHCP packet received on an untrusted port where the DHCP client hardware address field does not match the source MAC address in the packet	N/A
Unless configured to not perform this check, a DHCP packet containing DHCP relay information (option 82) received from an untrusted port	N/A
A broadcast packet that has a MAC address in the DHCP binding database, but the port in the DHCP binding database is different from the port on which the packet is received	DHCPRELEASE, DHCPDECLINE

Enabling DHCP Snooping

DHCP snooping is enabled globally by entering this command:

```
ProCurve(config)# dhcp-snooping
```

Use the **no** form of the command to disable DHCP snooping.

Syntax: [no] dhcp-snooping [authorized-server | database | option | trust | verify | vlan]

authorized server: Enter the IP address of a trusted DHCP server. If no authorized servers are configured, all DHCP server addresses are considered valid.
Maximum: 20 authorized servers

database: To configure a location for the lease database, enter a URL in the format **tftp://ip-addr/ascii-string**. The maximum number of characters for the URL is 63.

option: Add relay information option (Option 82) to DHCP client packets that are being forwarded out trusted ports. The default is **yes**, add relay information.

trust: Configure trusted ports. Only server packets received on trusted ports are forwarded. Default: **untrusted**.

verify: Enables DHCP packet validation. The DHCP client hardware address field and the source MAC address must be the same for packets received on untrusted ports or the packet is dropped. Default: **Yes**

vlan: Enable DHCP snooping on a vlan. DHCP snooping must be enabled already. Default: **No**

To display the DHCP snooping configuration, enter this command:

```
ProCurve(config)# show dhcp-snooping
```

An example of the output is shown below.

```
ProCurve(config)# show dhcp-snooping
DHCP Snooping Information
  DHCP Snooping           : Yes
  Enabled Vlans           :
  Verify MAC              : Yes
  Option 82 untrusted policy : drop
  Option 82 Insertion      : Yes
  Option 82 remote-id      : mac
  Store lease database     : Not configured
  Port  Trust
  ----  ----
  B1     No
  B2     No
```

Figure 8-1. An Example of the DHCP Snooping Command Output

To display statistics about the DHCP snooping process, enter this command:

```
ProCurve(config)# show dhcp-snooping stats
```

An example of the output is shown below.

```
ProCurve(config)# show dhcp-snooping stats
```

Packet type	Action	Reason	Count
server	forward	from trusted port	8
client	forward	to trusted port	8
server	drop	received on untrusted port	2
server	drop	unauthorized server	0
client	drop	destination on untrusted port	0
client	drop	untrusted option 82 field	0
client	drop	bad DHCP release request	0
client	drop	failed verify MAC check	0

Figure 8-2. Example of Show DHCP Snooping Statistics

Enabling DHCP Snooping on VLANs

DHCP snooping on VLANs is disabled by default. To enable DHCP snooping on a VLAN or range of VLANs enter this command:

```
ProCurve(config)# dhcp-snooping vlan <vlan-id-range>
```

You can also use this command in the vlan context, in which case you cannot enter a range of VLANs for snooping.

Below is an example of DHCP snooping enabled on VLAN 4.

```
ProCurve(config)# dhcp-snooping vlan 4
ProCurve(config)# show dhcp-snooping
```

DHCP Snooping Information

DHCP Snooping	: Yes
Enabled Vlans	: 4
Verify MAC	: Yes
Option 82 untrusted policy	: drop
Option 82 Insertion	: Yes
Option 82 remote-id	: mac

Figure 8-3. Example of DHCP Snooping on a VLAN

Configuring DHCP Snooping Trusted Ports

By default, all ports are untrusted. To configure a port or range of ports as trusted, enter this command:

```
ProCurve(config)# dhcp-snooping trust <port-list>
```

You can also use this command in the interface context, in which case you are not able to enter a list of ports.

```
ProCurve(config)# dhcp-snooping trust B1-B2
ProCurve(config)# show dhcp-snooping

DHCP Snooping Information

DHCP Snooping           : Yes
Enabled Vlans           : 4
Verify MAC              : Yes
Option 82 untrusted policy : drop
Option 82 Insertion      : Yes
Option 82 remote-id      : mac

Store lease database : Not configured

Port  Trust
-----
B1    Yes
B2    Yes
B3    No
```

Figure 8-4. Example of Setting Trusted Ports

DHCP server packets are forwarded only if received on a trusted port; DHCP server packets received on an untrusted port are dropped.

Use the **no** form of the command to remove the trusted configuration from a port.

Configuring Authorized Server Addresses

If authorized server addresses are configured, a packet from a DHCP server must be received on a trusted port AND have a source address in the authorized server list in order to be considered valid. If no authorized servers are configured, all servers are considered valid. You can configure a maximum of 20 authorized servers.

To configure a DHCP authorized server address, enter this command in the global configuration context:

```
ProCurve(config)# dhcp-snooping authorized-server  
                  <ip-address>
```

```
ProCurve(config)# show dhcp-snooping

DHCP Snooping Information

DHCP Snooping           : Yes
Enabled Vlans           : 4
Verify MAC              : No
Option 82 untrusted policy : drop
Option 82 Insertion     : Yes
Option 82 remote-id     : subnet-ip

Authorized Servers
-----
111.222.3.4
```

Figure 8-5. Example of Authorized Servers for DHCP Snooping

Using DHCP Snooping with Option 82

DHCP adds Option 82 (relay information option) to DHCP request packets received on untrusted ports by default. (See the preceding section *Configuring DHCP Relay* for more information on Option 82.)

When DHCP is enabled globally and also enabled on a VLAN, and the switch is acting as a DHCP relay, the settings for the DHCP relay Option 82 command are ignored when snooping is controlling Option 82 insertion. Option 82 inserted in this manner allows the association of the client's lease with the correct port, even when another device is acting as a DHCP relay or when the server is on the same subnet as the client.

Note

DHCP snooping only overrides the Option 82 settings on a VLAN that has snooping enabled, not on VLANs without snooping enabled.

If DHCP snooping is enabled on a switch where an edge switch is also using DHCP snooping, it is desirable to have the packets forwarded so the DHCP bindings are learned. To configure the policy for DHCP packets from untrusted ports that already have Option 82 present, enter this command in the global configuration context.

Syntax: [no] dhcp-snooping option 82 [remote-id <mac | subnet-ip | mgmt-ip>]
[untrusted-policy <drop | keep | replace>]

Enables DHCP Option 82 insertion in the packet.

remote-id *Set the value used for the **remote-id** field of the relay information option.*

mac: *The switch mac address is used for the remote-id. This is the default.*

subnet-ip: *The IP address of the VLAN the packet was received on is used for the remote-id. If **subnet-ip** is specified but the value is not set, the MAC address is used.*

mgmt-ip: *The management VLAN IP address is used as the remote-id. If **mgmt-ip** is specified but the value is not set, the MAC address is used.*

untrusted-policy *Configures DHCP snooping behavior when forwarding a DHCP packet from an untrusted port that already contains DHCP relay information (Option 82). The default is **drop**.*

drop: *The packet is dropped.*

keep: *The packet is forwarded without replacing the option information.*

replace: *The existing option is replaced with a new Option 82 generated by the switch.*

Note

The default **drop** policy should remain in effect if there are any untrusted nodes, such as clients, directly connected to this switch.

Changing the Remote-id from a MAC to an IP Address

By default, DHCP snooping uses the MAC address of the switch as the remote-id in Option 82 additions. The IP address of the VLAN the packet was received on or the IP address of the management VLAN can be used instead by entering this command with the associated parameter:

```
ProCurve(config)# dhcp-snooping option 82 remote-id  
                  <mac|subnet-ip|mgmt-ip>
```

```
ProCurve(config)# dhcp-snooping option 82 remote-id subnet-  
ip  
ProCurve(config)# show dhcp-snooping  
  
DHCP Snooping Information  
  
DHCP Snooping           : Yes  
Enabled Vlans           : 4  
Verify MAC              : Yes  
Option 82 untrusted policy : drop  
Option 82 Insertion      : Yes  
Option 82 remote-id      : subnet-ip
```

Figure 8-6. Example of DHCP Snooping Option 82 using the VLAN IP Address

Disabling the MAC Address Check

DHCP snooping drops DHCP packets received on untrusted ports when the check address (chaddr) field in the DHCP header does not match the source MAC address of the packet (default behavior). To disable this checking, use the **no** form of this command.

```
ProCurve(config)# dhcp-snooping verify mac
```

```
ProCurve(config)# dhcp-snooping verify mac
ProCurve(config)# show dhcp-snooping

DHCP Snooping Information

DHCP Snooping           : Yes
Enabled Vlans           : 4
Verify MAC              : yes
Option 82 untrusted policy : drop
Option 82 Insertion     : Yes
Option 82 remote-id     : subnet-ip
```

Figure 8-7. Example Showing the DHCP Snooping Verify MAC Setting

The DHCP Binding Database

DHCP snooping maintains a database of up to 8192 DHCP bindings on untrusted ports. Each binding consists of:

- Client MAC address
- Port number
- VLAN identifier
- Leased IP address
- Lease time

The switch can be configured to store the bindings at a specific URL so they will not be lost if the switch is rebooted. If the switch is rebooted, it will read its binding database from the specified location. To configure this location use this command.

Syntax: [no] dhcp-snooping database [file<tftp://<ip-address>/<ascii-string>>]
[delay<15-86400>][timeout<0-86400>]

file *Must be in Uniform Resource Locator (URL) format — “tftp://ip-address/ascii-string”. The maximum filename length is 63 characters.*

delay *Number of seconds to wait before writing to the database. Default = 300 seconds.*

timeout *Number of seconds to wait for the database file transfer to finish before returning an error. A value of zero (0) means retry indefinitely. Default = 300 seconds.*

A message is logged in the system event log if the DHCP binding database fails to update.

To display the contents of the DHCP snooping binding database, enter this command.

Syntax: show dhcp-snooping binding

```
ProCurve(config)# show dhcp-snooping binding
```

MacAddress	IP	VLAN	Interface	Time left
-----	-----	----	-----	-----
22.22.22.22.22.22	10.0.0.1	4	B2	1600

Figure 8-8. Example Showing DHCP Snooping Binding Database Contents

Note

If a lease database is configured, the switch drops all DHCP packets until the lease database is read. This only occurs when the switch reboots and is completed quickly. If the switch is unable to read the lease database from the tftp server, it waits until that operation times out and then begins forwarding DHCP packets.

Enabling Debug Logging

To enable debug logging for DHCP snooping, use this command.

Syntax: [no] debug dhcp-snooping [agent | event | packet]

agent *Displays DHCP snooping agent messages.*

event *Displays DHCP snooping event messages.*

packet *Displays DHCP snooping packet messages.*

Operational Notes

- DHCP is not configurable from the web management interface or menu interface.
- If packets are received at too high a rate, some may be dropped and need to be re-transmitted.

- ProCurve recommends running a time synchronization protocol such as SNTP in order to track lease times accurately.
- A remote server must be used to save lease information or there may be a loss of connectivity after a switch reboot.

Log Messages

Server <ip-address> packet received on untrusted port <port-number> dropped. Indicates a DHCP server on an untrusted port is attempting to transmit a packet. This event is recognized by the reception of a DHCP server packet on a port that is configured as untrusted.

Ceasing untrusted server logs for %s. More than one packet was received from a DHCP server on an untrusted port. To avoid filling the log file with repeated attempts, untrusted server drop packet events will not be logged for the specified <duration>.

Client packet destined to untrusted port <port-number> dropped. Indicates that the destination of a DHCP client unicast packet is on an untrusted port. This event is recognized when a client unicast packet is dropped because the destination address is out a port configured as untrusted.

Ceasing untrusted port destination logs for %s. More than one client unicast packet with an untrusted port destination was dropped. To avoid filling the log file with repeated attempts, untrusted port destination attempts will not be logged for the specified <duration>.

Unauthorized server <ip-address> detected on port <port-number>. Indicates that an unauthorized DHCP server is attempting to send packets. This event is recognized when a server packet is dropped because there are configured authorized servers and a server packet is received from a server that is not configured as an authorized server.

Ceasing unauthorized server logs for <duration>. More than one unauthorized server packet was dropped. To avoid filling the log file with repeated attempts, unauthorized server transmit attempts will not be logged for the specified <duration>.

Received untrusted relay information from client <mac-address> on port <port-number>. Indicates the reception on an untrusted port of a client packet containing a relay information option field. This event is recognized when a client packet containing a relay information option field is dropped because it was received on a port configured as untrusted.

Ceasing untrusted relay information logs for <duration>. More than one DHCP client packet received on an untrusted port with a relay information field was dropped. To avoid filling the log file with repeated attempts, untrusted relay information packets will not be logged for the specified <duration>.

Client address <mac-address> not equal to source MAC <mac-address> detected on port <port-number>. Indicates that a client packet source MAC address does not match the “chaddr” field. This event is recognized when the dhcp-snooping agent is enabled to filter DHCP client packets that do not have a matching “chaddr” field and source MAC address.

Ceasing MAC mismatch logs for <duration>. More than one DHCP client packet with a mismatched source MAC and chaddr field was dropped. To avoid filling the log file with repeated attempts, client address mismatch events will not be logged for the specified <duration>.

Attempt to release address <ip-address> leased to port <port-number> detected on port <port-number> dropped. Indicates an attempt by a client to release an address when a DHCPRELEASE or DHCPDECLINE packet is received on a port different from the port the address was leased to.

Ceasing bad release logs for %s. More than one bad DHCP client release packet was dropped. To avoid filling the log file with repeated bad release dropped packets, bad releases will not be logged for <duration>.

Lease table is full, DHCP lease was not added. The lease table is full and this lease will not be added to it.

Write database to remote file failed errno (error-num). An error occurred while writing the temporary file and sending it using tftp to the remote server.

DHCP packets being rate-limited. Too many DHCP packets are flowing through the switch and some are being dropped.

Snooping table is full. The DHCP binding table is full and subsequent bindings are being dropped.

Dynamic ARP Protection

Introduction

On the VLAN interfaces of a routing switch, dynamic ARP protection ensures that only valid ARP requests and responses are relayed or used to update the local ARP cache. ARP packets with invalid IP-to-MAC address bindings advertised in the source protocol address and source physical address fields are discarded. For more information about the ARP cache, refer to “ARP Cache Table” in the *Multicast and Routing Guide*.

ARP requests are ordinarily broadcast and received by all devices in a broadcast domain. Most ARP devices update their IP-to-MAC address entries each time they receive an ARP packet even if they did not request the information. This behavior makes an ARP cache vulnerable to attacks.

Because ARP allows a node to update its cache entries on other systems by broadcasting or unicasting a gratuitous ARP reply, an attacker can send his own IP-to-MAC address binding in the reply that causes all traffic destined for a VLAN node to be sent to the attacker's MAC address. As a result, the attacker can intercept traffic for other hosts in a classic "man-in-the-middle" attack. The attacker gains access to any traffic sent to the poisoned address and can capture passwords, e-mail, and VoIP calls or even modify traffic before resending it.

Another way in which the ARP cache of known IP addresses and associated MAC addresses can be poisoned is through unsolicited ARP responses. For example, an attacker can associate the IP address of the network gateway with the MAC address of a network node. In this way, all outgoing traffic is prevented from leaving the network because the node does not have access to outside networks. As a result, the node is overwhelmed by outgoing traffic destined to another network.

Dynamic ARP protection is designed to protect your network against ARP poisoning attacks in the following ways:

- Allows you to differentiate between trusted and untrusted ports.
- Intercepts all ARP requests and responses on untrusted ports before forwarding them.
- Verifies IP-to-MAC address bindings on untrusted ports with the information stored in the lease database maintained by DHCP snooping and user-configured static bindings (in non-DHCP environments):

- If a binding is valid, the switch updates its local ARP cache and forwards the packet.
- If a binding is invalid, the switch drops the packet, preventing other network devices from receiving the invalid IP-to-MAC information.

DHCP snooping intercepts and examines DHCP packets received on switch ports before forwarding the packets. DHCP packets are checked against a database of DHCP binding information. Each binding consists of a client MAC address, port number, VLAN identifier, leased IP address, and lease time. The DHCP binding database is used to validate packets by other security features on the switch. For more information, refer to “DHCP Snooping” in the *Access Security Guide*.

If you have already enabled DHCP snooping on a switch, you may also want to add static IP-to-MAC address bindings to the DHCP snooping database so that ARP packets from devices that have been assigned static IP addresses are also verified.

- Supports additional checks to verify source MAC address, destination MAC address, and IP address.

ARP packets that contain invalid IP addresses or MAC addresses in their body that do not match the addresses in the Ethernet header are dropped.

When dynamic ARP protection is enabled, only ARP request and reply packets with valid IP-to-MAC address bindings in their packet header are relayed and used to update the ARP cache.

Dynamic ARP protection is implemented in the following ways on a switch:

- You can configure dynamic ARP protection only from the CLI; you cannot configure this feature from the web or menu interfaces.
- Line rate—Dynamic ARP protection copies ARP packets to the switch CPU, evaluates the packets, and then re-forwards them through the switch software. During this process, if ARP packets are received at too high a line rate, some ARP packets may be dropped and will need to be retransmitted.
- The SNMP MIB, HP-ICF-ARP-PROTECT-MIB, is created to configure dynamic ARP protection and to report ARP packet-forwarding status and counters.

Enabling Dynamic ARP Protection

To enable dynamic ARP protection for VLAN traffic on a routing switch, enter the **arp-protect vlan** command at the global configuration level.

Syntax: [no] arp-protect vlan [*vlan-range*]

vlan-range *Specifies a VLAN ID or a range of VLAN IDs from one to 4094; for example, 1–200.*

An example of the **arp-protect vlan** command is shown here:

```
ProCurve(config)# arp-protect vlan 1-101
```

Configuring Trusted Ports

In a similar way to DHCP snooping, dynamic ARP protection allows you to configure VLAN interfaces in two categories: trusted and untrusted ports. ARP packets received on trusted ports are forwarded without validation.

By default, all ports on a switch are untrusted. If a VLAN interface is untrusted:

- The switch intercepts all ARP requests and responses on the port.
- Each intercepted packet is checked to see if its IP-to-MAC binding is valid. If a binding is invalid, the switch drops the packet.

You must configure trusted ports carefully. For example, in the topology in Figure 8-9, Switch B may not see the leased IP address that Host 1 receives from the DHCP server. If the port on Switch B that is connected to Switch A is untrusted and if Switch B has dynamic ARP protection enabled, it will see ARP packets from Host 1 as invalid, resulting in a loss of connectivity.

On the other hand, if Switch A does not support dynamic ARP protection and you configure the port on Switch B connected to Switch A as trusted, Switch B opens itself to possible ARP poisoning from hosts attached to Switch A.




Figure 8-9. Configuring Trusted Ports for Dynamic ARP Protection

Take into account the following configuration guidelines when you use dynamic ARP protection in your network:

- You should configure ports connected to other switches in the network as trusted ports. In this way, all network switches can exchange ARP packets and update their ARP caches with valid information.
- Switches that do not support dynamic ARP protection should be separated by a router in their own Layer 2 domain. Because ARP packets do not cross Layer 2 domains, the unprotected switches cannot unknowingly accept ARP packets from an attacker and forward them to protected switches through trusted ports.

To configure one or more Ethernet interfaces that handle VLAN traffic as trusted ports, enter the **arp-protect trust** command at the global configuration level. The switch does not check ARP requests and responses received on a trusted port.

Syntax: [no] arp-protect trust <port-list>

port-list	<i>Specifies a port number or a range of port numbers. Separate individual port numbers or ranges of port numbers with a comma; for example: c1-c3, c6.</i>
------------------	--

An example of the **arp-protect trust** command is shown here:

```
ProCurve(config)# arp-protect trust b1-b4, d1
```

Adding an IP-to-MAC Binding to the DHCP Database

A routing switch maintains a DHCP binding database, which is used for DHCP and ARP packet validation. Both the DHCP snooping and DHCP Option 82 insertion features maintain the lease database by learning the IP-to-MAC bindings on untrusted ports. Each binding consists of the client MAC address, port number, VLAN identifier, leased IP address, and lease time.

If your network does not use DHCP or if some network devices have fixed, user-configured IP addresses, you can enter static IP-to-MAC address bindings in the DHCP binding database. The switch uses manually configured static bindings for DHCP snooping and dynamic ARP protection.

To add the static configuration of an IP-to-MAC binding for a port to the database, enter the **ip source-binding** command at the global configuration level.

Syntax: [no] ip source-binding <mac-address> vlan <vlan-id> <ip-address>
interface <port-number>

mac-address *Specifies a MAC address to bind with a VLAN and IP address on the specified port in the DHCP binding database.*

vlan <vlan-id> *Specifies a VLAN ID number to bind with the specified MAC and IP addresses on the specified port in the DHCP binding database.*

ip-address *Specifies an IP address to bind with a VLAN and MAC address on the specified port in the DHCP binding database.*

interface <port-number> *Specifies the port number on which the IP-to-MAC address and VLAN binding is configured in the DHCP binding database.*

An example of the **ip source-binding** command is shown here:

```
ProCurve(config)# ip source-binding 0030c1-7f49c0  
interface vlan 100 10.10.20.1 interface A4
```

Note

Note that the **ip source-binding** command is the same command used by the Dynamic IP Lockdown feature to configure static bindings. The Dynamic ARP Protection and Dynamic IP Lockdown features share a common list of source IP-to-MAC bindings.

Configuring Additional Validation Checks on ARP Packets

Dynamic ARP protection can be configured to perform additional validation checks on ARP packets. By default, no additional checks are performed. To configure additional validation checks, enter the **arp-protect validate** command at the global configuration level.

Syntax: [no] arp-protect validate <[src-mac] | [dst-mac] | [ip]>

- | | |
|----------------|--|
| src-mac | <i>(Optional) Drops any ARP request or response packet in which the source MAC address in the Ethernet header does not match the sender MAC address in the body of the ARP packet.</i> |
| dst-mac | <i>(Optional) Drops any unicast ARP response packet in which the destination MAC address in the Ethernet header does not match the target MAC address in the body of the ARP packet.</i> |
| ip | <i>(Optional) Drops any ARP packet in which the sender IP address is invalid. Drops any ARP response packet in which the target IP address is invalid. Invalid IP addresses include: 0.0.0.0, 255.255.255.255, all IP multicast addresses, and all Class E IP addresses.</i> |

You can configure one or more of the validation checks. The following example of the **arp-protect validate** command shows how to configure the validation checks for source MAC address and destination MAC address:

```
ProCurve(config)# arp-protect validate src-mac dst-mac
```

Verifying the Configuration of Dynamic ARP Protection

To display the current configuration of dynamic ARP protection, including the additional validation checks and the trusted ports that are configured, enter the **show arp-protect** command:

```
ProCurve(config)# show arp-protect

ARP Protection Information

Enabled Vlans   : 1-4094
Validate       : dst-mac, src-mac

Port   Trust
-----
B1     Yes
B2     Yes
B3     No
B4     No
B5     No
```

Figure 8-1. The show arp-protect Command

Displaying ARP Packet Statistics

To display statistics about forwarded ARP packets, dropped ARP packets, MAC validation failure, and IP validation failures, enter the **show arp-protect statistics** command:

```
ProCurve(config)# show arp-protect statistics

Status and Counters - ARP Protection Counters for VLAN 1

Forwarded pkts   : 10      Bad source mac       : 2
Bad bindings     : 1       Bad destination mac: 1
Malformed pkts  : 0       Bad IP address       : 0

Status and Counters - ARP Protection Counters for VLAN 2

Forwarded pkts   : 1       Bad source mac       : 1
Bad bindings     : 1       Bad destination mac: 1
Malformed pkts  : 1       Bad IP address       : 1
```

Figure 8-2. Show arp-protect statistics Command

Monitoring Dynamic ARP Protection

When dynamic ARP protection is enabled, you can monitor and troubleshoot the validation of ARP packets with the **debug arp-protect** command. Use this command when you want to debug the following conditions:

- The switch is dropping valid ARP packets that should be allowed.
- The switch is allowing invalid ARP packets that should be dropped.

```
ProCurve(config)# debug arp-protect

1. ARP request is valid
"DARPP: Allow ARP request 000000-000001,10.0.0.1 for 10.0.0.2 port A1,
vlan "

2. ARP request detected with an invalid binding
"DARPP: Deny ARP request 000000-000003,10.0.0.1 port A1, vlan 1"

3. ARP response with a valid binding
"DARPP: Allow ARP reply 000000-000002,10.0.0.2 port A2, vlan 1"

4. ARP response detected with an invalid binding
"DARPP: Deny ARP reply 000000-000003,10.0.0.2 port A2, vlan 1"
```

Figure 8-3. Example of debug arp-protect Command

Dynamic IP Lockdown

The Dynamic IP Lockdown feature is used to prevent IP source address spoofing on a per-port and per-VLAN basis. When dynamic IP lockdown is enabled, IP packets in VLAN traffic received on a port are forwarded only if they contain a known source IP address and MAC address binding for the port. The IP-to-MAC address binding can either be statically configured or learned by the DHCP Snooping feature.

Protection Against IP Source Address Spoofing

Many network attacks occur when an attacker injects packets with forged IP source addresses into the network. Also, some network services use the IP source address as a component in their authentication schemes. For example, the BSD “r” protocols (rlogin, rcp, rsh) rely on the IP source address for packet authentication. SNMPv1 and SNMPv2c also frequently use authorized IP address lists to limit management access. An attacker that is able to send traffic that appears to originate from an authorized IP source address may gain access to network services for which he is not authorized.

Dynamic IP lockdown provides protection against IP source address spoofing by means of IP-level port security. IP packets received on a port enabled for dynamic IP lockdown are only forwarded if they contain a known IP source address and MAC address binding for the port.

Dynamic IP lockdown uses information collected in the DHCP Snooping lease database and through statically configured IP source bindings to create internal, per-port lists. The internal lists are dynamically created from known IP-to-MAC address bindings to filter VLAN traffic on both the source IP address and source MAC address.

Prerequisite: DHCP Snooping

Dynamic IP lockdown requires that you enable DHCP snooping as a prerequisite for its operation on ports and VLAN traffic:

- Dynamic IP lockdown only enables traffic for clients whose leased IP addresses are already stored in the lease database created by DHCP snooping or added through a static configuration of an IP-to-MAC binding.

Therefore, if you enable DHCP snooping after dynamic IP lockdown is enabled, clients with an existing DHCP-assigned address must either request a new leased IP address or renew their existing DHCP-assigned address. Otherwise, a client's leased IP address is not contained in the DHCP binding database. As a result, dynamic IP lockdown will not allow inbound traffic from the client.

- It is recommended that you enable DHCP snooping a week before you enable dynamic IP lockdown to allow the DHCP binding database to learn clients' leased IP addresses. You must also ensure that the lease time for the information in the DHCP binding database lasts more than a week.

Alternatively, you can configure a DHCP server to re-allocate IP addresses to DHCP clients. In this way, you repopulate the lease database with current IP-to-MAC bindings.

- The DHCP binding database allows VLANs enabled for DHCP snooping to be known on ports configured for dynamic IP lockdown. As new IP-to-MAC address and VLAN bindings are learned, a corresponding permit rule is dynamically created and applied to the port (preceding the final deny any vlan <VLAN_IDs> rule as shown in the example in Figure 8-4. These VLAN_IDs correspond to the subset of configured and enabled VLANs for which DHCP snooping has been configured.
- For dynamic IP lockdown to work, a port must be a member of at least one VLAN that has DHCP snooping enabled.
- Disabling DHCP snooping on a VLAN causes Dynamic IP bindings on Dynamic IP Lockdown-enabled ports in this VLAN to be removed. The port reverts back to switching traffic as usual.

Filtering IP and MAC Addresses Per-Port and Per-VLAN

This section contains an example that shows the following aspects of the Dynamic IP Lockdown feature:

- Internal Dynamic IP lockdown bindings dynamically applied on a per-port basis from information in the DHCP Snooping lease database and statically configured IP-to-MAC address bindings
- Packet filtering using source IP address, source MAC address, and source VLAN as criteria

In this example, the following DHCP leases have been learned by DHCP snooping on port 5. VLANs 2 and 5 are enabled for DHCP snooping.

Table 1. Sample DHCP Snooping Entries

IP Address	MAC Address	VLAN ID
10.0.8.5	001122-334455	2
10.0.8.7	001122-334477	2
10.0.10.3	001122-334433	5

The following example shows an IP-to-MAC address and VLAN binding that have been statically configured in the lease database on port 5.

IP Address	MAC Address	VLAN ID
10.0.10.1	001122-110011	5

Assuming that DHCP snooping is enabled and that port 5 is untrusted, dynamic IP lockdown applies the following dynamic VLAN filtering on port 5:

```
permit 10.0.8.5 001122-334455 vlan 2  
  
permit 10.0.8.7 001122-334477 vlan 2  
  
permit 10.0.10.3 001122-334433 vlan 5  
  
permit 10.0.10.1 001122-110011 vlan 5  
  
deny any vlan 1-10  
  
permit any
```

Figure 8-4. Example of Internal Statements used by Dynamic IP Lockdown

Note that the **deny any** statement is applied only to VLANs for which DHCP snooping is enabled. The **permit any** statement is applied only to all other VLANs.

Enabling Dynamic IP Lockdown

To enable dynamic IP lockdown on all ports or specified ports, enter the **ip source-lockdown** command at the global configuration level. Use the **no** form of the command to disable dynamic IP lockdown.

Syntax: [no] ip source-lockdown <port-list>

Enables dynamic IP lockdown globally on all ports or on specified ports on the routing switch.

Operating Notes

- Dynamic IP lockdown is enabled at the port configuration level and applies to all bridged or routed IP packets entering the switch. The only IP packets that are exempt from dynamic IP lockdown are broadcast DHCP request packets, which are handled by DHCP snooping.
- DHCP snooping is a prerequisite for Dynamic IP Lockdown operation. The following restrictions apply:
 - DHCP snooping is required for dynamic IP lockdown to operate. To enable DHCP snooping, enter the **dhcp-snooping** command at the global configuration level.

- Dynamic IP lockdown only filters packets in VLANs that are enabled for DHCP snooping. In order for Dynamic IP lockdown to work on a port, the port must be configured for at least one VLAN that is enabled for DHCP snooping.

To enable DHCP snooping on a VLAN, enter the **dhcp-snooping vlan** [*vlan-id-range*] command at the global configuration level or the **dhcp-snooping** command at the VLAN configuration level.

- Dynamic IP lockdown is not supported on a trusted port. (However, note that the DHCP server must be connected to a trusted port when DHCP snooping is enabled.)

By default, all ports are untrusted. To remove the trusted configuration from a port, enter the **no dhcp-snooping trust** <*port-list*> command at the global configuration level.

For more information on how to configure and use DHCP snooping, see “DHCP Snooping” on page 8-4.

- After you enter the **ip source-lockdown** command (enabled globally with the desired ports entered in <*port-list*>), the dynamic IP lockdown feature remains disabled on a port if any of the following conditions exist:
 - If DHCP snooping has not been globally enabled on the switch.
 - If the port is not a member of at least one VLAN that is enabled for DHCP snooping.
 - If the port is configured as a trusted port for DHCP snooping.

Dynamic IP lockdown is activated on the port only after you make the following configuration changes:

- Enable DHCP snooping on the switch.
- Configure the port as a member of a VLAN that has DHCP snooping enabled.
- Remove the trusted-port configuration.
- You can configure dynamic IP lockdown only from the CLI; this feature cannot be configured from the web management or menu interface.
- If you enable dynamic IP lockdown on a port, you cannot add the port to a trunk.
- Dynamic IP lockdown must be removed from a trunk before the trunk is removed.

Adding an IP-to-MAC Binding to the DHCP Binding Database

A switch maintains a DHCP binding database, which is used for dynamic IP lockdown as well as for DHCP and ARP packet validation. The DHCP snooping feature maintains the lease database by learning the IP-to-MAC bindings of VLAN traffic on untrusted ports. Each binding consists of the client MAC address, port number, VLAN identifier, leased IP address, and lease time.

Dynamic IP lockdown supports a total of 4K static and dynamic bindings with up to 64 bindings per port. When DHCP snooping is enabled globally on a VLAN, dynamic bindings are learned when a client on the VLAN obtains an IP address from a DHCP server. Static bindings are created manually with the CLI or from a downloaded configuration file.

When dynamic IP lockdown is enabled globally or on ports the bindings associated with the ports are written to hardware. This occurs during these events:

- Switch initialization
- Hot swap
- A dynamic IP lockdown-enabled port is moved to a DHCP snooping-enabled VLAN
- DHCP snooping or dynamic IP lockdown characteristics are changed such that dynamic IP lockdown is enabled on the ports

Potential Issues with Bindings

- When dynamic IP lockdown enabled, and a port or switch has the maximum number of bindings configured, the client DHCP request will be dropped and the client will not receive an IP address through DHCP.
- When dynamic IP lockdown is enabled and a port is configured with the maximum number of bindings, adding a static binding to the port will fail.
- When dynamic IP lockdown is enabled globally, the bindings for each port are written to hardware. If global dynamic IP lockdown is enabled and disabled several times, it is possible to run out of buffer space for additional bindings. The software will delay adding the bindings to hardware until resources are available.

Adding a Static Binding

To add the static configuration of an IP-to-MAC binding for a port to the lease database, enter the **ip source-binding** command at the global configuration level. Use the **no** form of the command to remove the IP-to-MAC binding from the database.

Syntax:	[no] ip source-binding <vlan-id> <ip-address> <mac-address> <port-number>
vlan-id	Specifies a valid VLAN ID number to bind with the specified MAC and IP addresses on the port in the DHCP binding database.
ip-address	Specifies a valid client IP address to bind with a VLAN and MAC address on the port in the DHCP binding database.
mac-address	Specifies a valid client MAC address to bind with a VLAN and IP address on the port in the DHCP binding database.
port-number	Specifies the port number on which the IP-to-MAC address and VLAN binding is configured in the DHCP binding database.

Note	Note that the ip source-binding command is the same command used by the Dynamic ARP Protection feature to configure static bindings. The Dynamic ARP Protection and Dynamic IP Lockdown features share a common list of source IP-to-MAC address bindings.
-------------	---

Verifying the Dynamic IP Lockdown Configuration

To display the ports on which dynamic IP lockdown is configured, enter the **show ip source-lockdown status** command at the global configuration level.

Syntax: show ip source-lockdown status

An example of the **show ip source-lockdown status** command output is shown in Figure 8-5. Note that the operational status of all switch ports is displayed. This information indicates whether or not dynamic IP lockdown is supported on a port.

```
ProCurve(config)# show ip source-lockdown status
Dynamic IP Lockdown (DIPLD) Information

Global State: Enabled

      Port      Operational State
      - - - - -  - - - - -
      A1         Active
      A2         Not in DHCP Snooping vlan
      A3         Disabled
      A4         Disabled
      A5         Trusted port, Not in DHCP Snooping vlan
      . . . . .
```

Figure 8-5. Example of show ip source-lockdown status Command Output

Displaying the Static Configuration of IP-to-MAC Bindings

To display the static configurations of IP-to-MAC bindings stored in the DHCP lease database, enter the **show ip source-lockdown bindings** command.

Syntax: show ip source-lockdown bindings [*port-number*]

port-number (Optional) Specifies the port number on which source IP-to-MAC address and VLAN bindings are configured in the DHCP lease database.

An example of the **show ip source-lockdown bindings** command output is shown in Figure 8-6.

```
ProCurve(config)# show ip source-lockdown bindings
```

Dynamic IP Lockdown (DIPLD) Bindings				
Mac Address	IP Address	VLAN	Port	Not in HW
-----	-----	-----	-----	-----
001122-334455	10.10.10.1	1111	X11	
005544-332211	10.10.10.2	2222	Trk11	YES
.

Figure 8-6. Example of show ip source-lockdown bindings Command Output

In the **show ip source-lockdown bindings** command output, the “Not in HW” column specifies whether or not (YES or NO) a statically configured IP-to-MAC and VLAN binding on a specified port has been combined in the lease database maintained by the DHCP Snooping feature.

Debugging Dynamic IP Lockdown

To enable the debugging of packets dropped by dynamic IP lockdown, enter the **debug dynamic-ip-lockdown** command.

Syntax: debug dynamic-ip-lockdown

To send command output to the active CLI session, enter the **debug destination session** command.

Counters for denied packets are displayed in the **debug dynamic-ip-lockdown** command output. Packet counts are updated every five minutes. An example of the command output is shown in Figure 8-7.

When dynamic IP lockdown drops IP packets in VLAN traffic that do not contain a known source IP-to-MAC address binding for the port on which the packets are received, a message is entered in the event log.


```
ProCurve(config)# debug dynamic-ip-lockdown

DIPLD 01/01/90 00:01:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 1 packets
DIPLD 01/01/90 00:06:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 294 packets
DIPLD 01/01/90 00:11:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
DIPLD 01/01/90 00:16:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
DIPLD 01/01/90 00:21:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 299 packets
DIPLD 01/01/90 00:26:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
DIPLD 01/01/90 00:31:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
DIPLD 01/01/90 00:36:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 299 packets
DIPLD 01/01/90 00:41:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
```

Figure 8-7. Example of debug dynamic-ip-lockdown Command Output

Using the Instrumentation Monitor

The instrumentation monitor can be used to detect anomalies caused by security attacks or other irregular operations on the switch. The following table shows the operating parameters that can be monitored at pre-determined intervals, and the possible security attacks that may trigger an alert:

Parameter Name	Description
pkts-to-closed-ports	The count of packets per minute sent to closed TCP/UDP ports. An excessive amount of packets could indicate a port scan, in which an attacker is attempting to expose a vulnerability in the switch.
arp-requests	The count of ARP requests processed per minute. A large amount of ARP request packets could indicate an host infected with a virus that is trying to spread itself.
ip-address-count	The number of destination IP addresses learned in the IP forwarding table. Some attacks fill the IP forwarding table causing legitimate traffic to be dropped.
system-resource-usage	The percentage of system resources in use. Some Denial-of-Service (DoS) attacks will cause excessive system resource usage, resulting in insufficient resources for legitimate traffic.
login-failures/min	The count of failed CLI login attempts or SNMP management authentication failures. This indicates an attempt has been made to manage the switch with an invalid login or password. Also, it might indicate a network management station has not been configured with the correct SNMP authentication parameters for the switch.
port-auth-failures/min	The count of times a client has been unsuccessful logging into the network
system-delay	The response time, in seconds, of the CPU to new network events such as BPDUs or packets for other network protocols. Some DoS attacks can cause the CPU to take too long to respond to new network events, which can lead to a breakdown of Spanning Tree or other features. A delay of several seconds indicates a problem.
mac-address-count	The number of MAC addresses learned in the forwarding table. Some attacks fill the forwarding table so that new conversations are flooded to all parts of the network.
mac-moves/min	The average number of MAC address moves from one port to another per minute. This usually indicates a network loop, but can also be caused by DoS attacks.
learn-discards/min	Number of MAC address learn events per minute discarded to help free CPU resources when busy.

Operating Notes

- To generate alerts for monitored events, you must enable the instrumentation monitoring log and/or SNMP trap. The threshold for each monitored parameter can be adjusted to minimize false alarms (see “Configuring Instrumentation Monitor” on page 8-35).
- When a parameter exceeds its threshold, an alert (event log message and/or SNMP trap) is generated to inform network administrators of this condition. The following example shows an event log message that occurs when the number of MAC addresses learned in the forwarding table exceeds the configured threshold:




Figure 8-8. Example of Event Log Message generated by Instrumentation Monitor

- Alerts are automatically rate limited to prevent filling the log file with redundant information. The following is an example of alerts that occur when the device is continually subject to the same attack (too many MAC addresses in this instance):

```
W 01/01/90 00:05:00 inst-mon: Limit for MAC addr count (300) is exceeded (321)
W 01/01/90 00:10:00 inst-mon: Limit for MAC addr count (300) is exceeded (323)
W 01/01/90 00:15:00 inst-mon: Limit for MAC addr count (300) is exceeded (322)
W 01/01/90 00:20:00 inst-mon: Limit for MAC addr count (300) is exceeded (324)
W 01/01/90 00:20:00 inst-mon: Ceasing logs for MAC addr count for 15 minutes
```

Figure 8-9. Example of rate limiting when multiple messages are generated

In the preceding example, if a condition is reported 4 times (persists for more than 15 minutes) then alerts cease for 15 minutes. If after 15 minutes the condition still exists, the alerts cease for 30 minutes, then for 1 hour, 2 hours, 4 hours, 8 hours, and after that the persisting condition is reported once a day. As with other event log entries, these alerts can be sent to a syslog server.

- **Known Limitations:** The instrumentation monitor runs once every five minutes. The current implementation does not track information such as the port, MAC, and IP address from which an attack is received.

Configuring Instrumentation Monitor

The following commands and parameters are used to configure the operational thresholds that are monitored on the switch. By default, the instrumentation monitor is disabled.

Syntax: [no] instrumentation monitor [parameterName|all] [<low|med|high|limitValue>]

[log] : Enables/disables instrumentation monitoring log so that event log messages are generated every time there is an event which exceeds a configured threshold.

(Default threshold setting when instrumentation monitoring is enabled: **enabled**)

[all] : Enables/disables all counter types on the switch but does not enable/disable instrumentation monitor logging.

(Default threshold setting when enabled: **see parameter listings below**)

[arp-requests] : The number of arp requests that are processed each minute.

(Default threshold setting when enabled: **1000 (med)**)

[ip-address-count]: The number of destination IP addresses learned in the IP forwarding table.

(Default threshold setting when enabled: **1000 (med)**)

[learn-discards]: The number of MAC address learn events per minute discarded to help free CPU resources when busy.

(Default threshold setting when enabled: **100 (med)**)

[login-failures]: The count of failed CLI login attempts or SNMP management authentication failures per hour.

(Default threshold setting when enabled: **10 (med)**)

[mac-address-count] : The number of MAC addresses learned in the forwarding table. You must enter a specific value in order to enable this feature.

(Default threshold setting when enabled: **1000 (med)**)

[mac-moves] : The average number of MAC address moves per minute from one port to another.

(Default threshold setting when enabled: **100 (med)**)

[pkts-to-closed-ports] : The count of packets per minute sent to closed TCP/UDP ports.

(Default threshold setting when enabled: **10 (med)**)

[port-auth-failures] : The count of times per minute that a client has been unsuccessful logging into the network.

(Default threshold setting when enabled: **10 (med)**)

[system-resource-usage]: The percentage of system resources in use.

(Default threshold setting when enabled: **50 (med)**)

[system-delay] : The response time, in seconds, of the CPU to new network events such as BPDU packets or packets for other network protocols.

(Default threshold setting when enabled: **3 seconds (med)**)

[trap] : Enables or disables SNMP trap generation.

(Default setting when instrumentation monitoring is enabled: **disabled**)

To enable instrumentation monitor using the default parameters and thresholds, enter the general **instrumentation monitor** command. To adjust specific settings, enter the name of the parameter that you wish to modify, and revise the threshold limits as needed.

Examples

To turn on monitoring and event log messaging with the default medium values:

```
ProCurve(config)# instrumentation monitor
```

To turn off monitoring of the system delay parameter:

```
ProCurve(config)# no instrumentation monitor system-  
delay
```

To adjust the alert threshold for the MAC address count to the low value:

```
ProCurve(config)# instrumentation monitor mac-  
address-count low
```

To adjust the alert threshold for the MAC address count to a specific value:

```
ProCurve(config)# instrumentation monitor mac-  
address-count 767
```

To enable monitoring of learn discards with the default medium threshold value:

```
ProCurve(config)# instrumentation monitor learn-  
discards
```

To disable monitoring of learn discards:

```
ProCurve(config)# no instrumentation monitor learn-  
discards
```

To enable or disable SNMP trap generation:

```
ProCurve(config)# [no] instrumentation monitor trap
```

Viewing the Current Instrumentation Monitor Configuration

The **show instrumentation monitor configuration** command displays the configured thresholds for monitored parameters.

```
ProCurve# show instrumentation monitor configuration
```

PARAMETER	LIMIT
mac-address-count	1000 (med)
ip-address-count	1000 (med)
system-resource-usage	50 (med)
system-delay	5 (high)
mac-moves/min	100 (med)
learn-discards/min	100 (med)
ip-port-scans/min	10 (med)
arp-requests/min	100 (low)
login-failures/min	10 (med)
port-auth-failures/min	10 (med)

SNMP trap generation for alerts: enabled
Instrumentation monitoring log : enabled

Figure 8-10. Viewing the Instrumentation Monitor Configuration

An alternate method of determining the current Instrumentation Monitor configuration is to use the **show run** command. However, the show run command output does not display the threshold values for each limit set.

Traffic/Security Filters and Monitors

Contents

Overview	9-2
Introduction	9-2
Filter Limits	9-2
Using Port Trunks with Filters	9-2
Filter Types and Operation	9-3
Source-Port Filters	9-4
Operating Rules for Source-Port Filters	9-4
Example	9-5
Named Source-Port Filters	9-6
Operating Rules for Named Source-Port Filters	9-6
Defining and Configuring Named Source-Port Filters	9-7
Viewing a Named Source-Port Filter	9-8
Using Named Source-Port Filters	9-9
Configuring Traffic/Security Filters	9-15
Configuring a Source-Port Traffic Filter	9-16
Example of Creating a Source-Port Filter	9-17
Configuring a Filter on a Port Trunk	9-17
Editing a Source-Port Filter	9-18
Filter Indexing	9-19
Displaying Traffic/Security Filters	9-20

Overview

Source-port filters are available on the HP ProCurve switch models covered in this guide.

Introduction

Feature	Default	Menu	CLI	Web
configure source-port filters	none	n/a	page 9-18	n/a
display filter data	n/a	n/a	page 9-20	n/a

You can enhance in-band security and improve control over access to network resources by configuring static filters to forward (the default action) or drop unwanted traffic. That is, you can configure a traffic filter to either forward or drop all network traffic moving to outbound (destination) ports and trunks (if any) on the switch.

Filter Limits

The switch accepts up to 101 static filters. These limitations apply:

- Source-port filters: up to 78

Using Port Trunks with Filters

The switch manages a port trunk as a single source or destination for source-port filtering. If you configure a port for filtering before adding it to a port trunk, the port retains the filter configuration, but suspends the filtering action while a member of the trunk. If you want a trunk to perform filtering, first configure the trunk, then configure the trunk for filtering. Refer to “Configuring a Filter on a Port Trunk” on page 9-17.

Filter Types and Operation

Table 9-1. Filter Types and Criteria

Static Filter Type	Selection Criteria
Source-Port	Inbound traffic from a designated, physical source-port will be forwarded or dropped on a per-port (destination) basis.

Source-Port Filters

This filter type enables the switch to forward or drop traffic from *all* end nodes on the indicated source-port to specific destination ports.

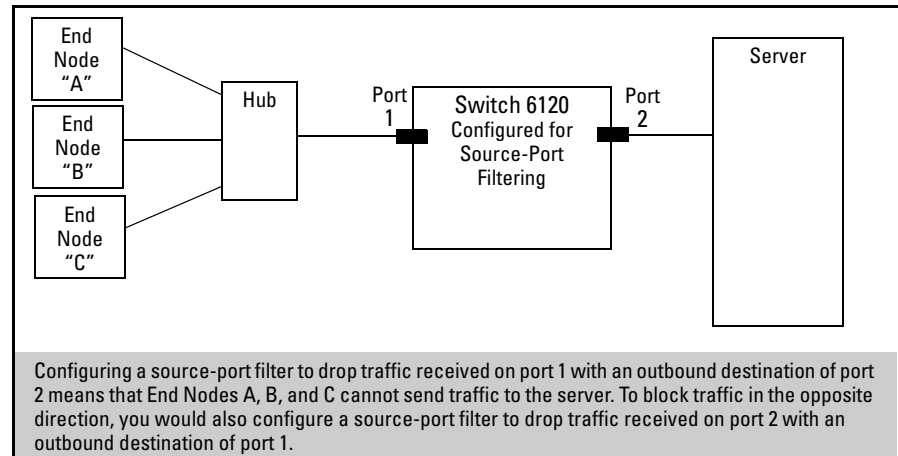


Figure 9-1. Example of a Source-Port Filter Application

Operating Rules for Source-Port Filters

- You can configure one source-port filter for each physical port and port trunk on the switch. (Refer to the **filter** command on page 9-16.)
- You can include all destination ports and trunks in the switch on a single source-port filter.
- Each source-port filter includes:
 - One source port or port trunk (**trk1**, **trk2**, ...**trkn**)
 - A set of destination ports and/or port trunks that includes all untrunked LAN ports and port trunks on the switch
 - An action (forward or drop) for each destination port or port trunk

When you create a source-port filter, the switch automatically sets the filter to forward traffic from the designated source to all destinations for which you do not specifically configure a “drop” action. Thus, it is not necessary to configure a source-port filter for traffic you want the switch to forward unless the filter was previously configured to drop the desired traffic.

- When you create a source port filter, all ports and port trunks (if any) on the switch appear as destinations on the list for that filter, even if routing is disabled and separate VLANs and/or subnets exist. Where traffic would normally be allowed between ports and/or trunks, the switch automatically forwards traffic to the outbound ports and/or trunks you do not specifically configure to drop traffic. (Destination ports that comprise a trunk are listed collectively by the trunk name—such as **Trk1**—instead of by individual port name.)
- Packets allowed for forwarding by a source-port filter are subject to the same operation as inbound packets on a port that is not configured for source-port filtering.
- With multiple IP addresses configured on a VLAN, and routing enabled on the switch, a single port or trunk can be both the source and destination of packets moving between subnets in that same VLAN. In this case, you can prevent the traffic of one subnet from being routed to another subnet of the same port by configuring the port or trunk as both the source and destination for traffic to drop.

Example

If you wanted to prevent server “A” from receiving traffic sent by workstation “X”, but do not want to prevent any other servers or end nodes from receiving traffic from workstation “X”, you would configure a filter to drop traffic from port 5 to port 7. The resulting filter would drop traffic from port 5 to port 7, but would forward all other traffic from any source port to any destination port. (Refer to figures 9-2 and 9-3.




Figure 9-2. Example of a Filter Blocking Traffic only from Port 5 to Server "A"

Traffic/Security Filters		
Filter Type : Source Port		
Source Port : 5		
Dest Port Type		Action
1	100/1000T	Forward
2	100/1000T	Forward
3	100/1000T	Forward
4	100/1000T	Forward
5	100/1000T	Forward
6	100/1000T	Forward
7	100/1000T	Drop
8	100/1000T	Forward
9	100/1000T	Forward
10	100/1000T	Forward
.	.	.
.	.	.
.	.	.
22	100/1000T	Forward
23	100/1000T	Forward
24	100/1000T	Forward

This list shows the filter created to block (drop) traffic from source port 5 (workstation "X") to destination port 7 (server "A"). Notice that the filter allows traffic to move from source port 5 to all other destination ports.

Figure 9-3. The Filter for the Actions Shown in Figure 9-2

Named Source-Port Filters

You can specify named source-port filters that may be used on multiple ports and port trunks. A port or port trunk can only have one source-port filter, but by using this capability you can define a source-port filter once and apply it to multiple ports and port trunks. This can make it easier to configure and manage source-port filters on your switch. The commands to define, configure, apply, and display the status of named source-port filters are described below.

Operating Rules for Named Source-Port Filters

- A port or port trunk may only have one source-port filter, named or not named.
- A named source-port filter can be applied to multiple ports or port trunks.
- Once a named source-port filter is defined, subsequent changes only modify its action, they don't replace it.
- To change the named source-port filter used on a port or port trunk, the current filter must first be removed, using the **no filter source-port named-filter <filter-name>** command.

- A named source-port filter can only be deleted when it is not applied to any ports.

Defining and Configuring Named Source-Port Filters

The named source-port filter command operates from the global configuration level.

Syntax: [no] filter source-port named-filter <filter-name>

Defines or deletes a named source-port filter. The <filter-name> may contain a maximum of 20 alpha-numeric characters (longer names may be specified, but they are not displayed.) A filter-name cannot be a valid port or port trunk name.

The maximum number of named source-port filters that can be used is equal to the number of ports on a switch.

*A named source-port filter can only be removed if it is not in use (use the **show filter source-port** command to check the status). Named source-port filters are not automatically deleted when they are no longer used.*

*Use the **no** option to delete an unused named source-port filter.*

Syntax: filter source-port named-filter <filter-name> drop < destination-port-list >

*Configures the named source-port filter to drop traffic having a destination on the ports and/or port trunks in the < destination-port-list >. Can be followed by the **forward** option if you have other destination ports or port trunks previously set to **drop** that you want to change to **forward**. For example:*

filter source-port named-filter <filter-name> drop < destination-port-list > forward < destination-port-list >

*The **destination-port-list** may contain ports, port trunks, and ranges (for example 3-7 or trk4-trk9) separated by commas.*

Syntax: filter source-port named-filter <filter-name> forward
< destination-port-list >

*Configures the named source-port filter to forward traffic having a destination on the ports and/or port trunks in the <**destination-port-list**>. Since “forward” is the default state for destinations in a filter, this command is useful when destinations in an existing filter are configured for “drop” and you want to change them to “forward”. Can be followed by the **drop** option if you have other destination ports set to **forward** that you want to change to **drop**. For example:*

filter source-port named-filter <filter-name> forward < destination-port-list > drop < destination-port-list >

A named source-port filter must first be defined and configured before it can be applied. In the following example two named source-port filters are defined, **web-only** and **accounting**.

```
ProCurve(config)# filter source-port named-filter web-only
ProCurve(config)# filter source-port named-filter accounting
```

By default, these two named source-port filters forward traffic to all ports and port trunks.

To configure a named source-port filter to prevent inbound traffic from being forwarded to specific destination switch ports or port trunks, the **drop** option is used. For example, on a 26-port switch, to configure the named source-port filter **web-only** to drop any traffic except that for destination ports 1 and 2, the following command would be used:

```
ProCurve(config)# filter source-port named-filter web-only drop 3-26
```

A named source-port filter can be defined and configured in a single command by adding the **drop** option, followed by the required destination-port-list.

Viewing a Named Source-Port Filter

You can list all source-port filters configured in the switch, both named and unnamed, and their action using the **show** command below.

Syntax: show filter source-port

Displays a listing of configured source-port filters, where each filter entry includes a Filter Name, Port List, and Action:

Filter Name: The filter-name used when a named source-port filter is defined. Non-named source-port filters are automatically assigned the port or port trunk number of the source port.

Port List: Lists the port and port trunk destinations using the filter. Named source-port filters that are not in use display **NOT USED**.

Action: Lists the ports and port trunks dropped by the filter. If a named source-port filter has been defined but not configured, this field is blank.

[index] For the supplied index (IDX) displays the action taken (Drop or Forward) for each destination port on the switch.

Using Named Source-Port Filters

A company wants to manage traffic to the Internet and its accounting server on a 26-port switch. Their network is pictured in Figure 9-4. Switch port 1 connects to a router that provides connectivity to a WAN and the Internet. Switch port 7 connects to the accounting server. Two workstations in accounting are connected to switch ports 10 and 11. Two workstations in engineering are connected to switch ports 12 and 13.

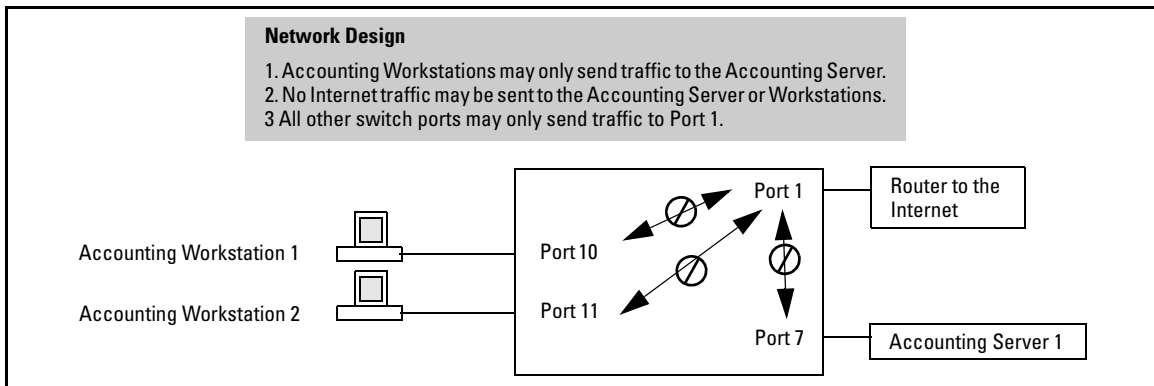


Figure 9-4. Network Configuration for Named Source-Port Filters Example

Defining and Configuring Example Named Source-Port Filters. While named source-port filters may be defined and configured in two steps, this is not necessary. Here we define and configure each of the named source-port filters for our example network in a single step.

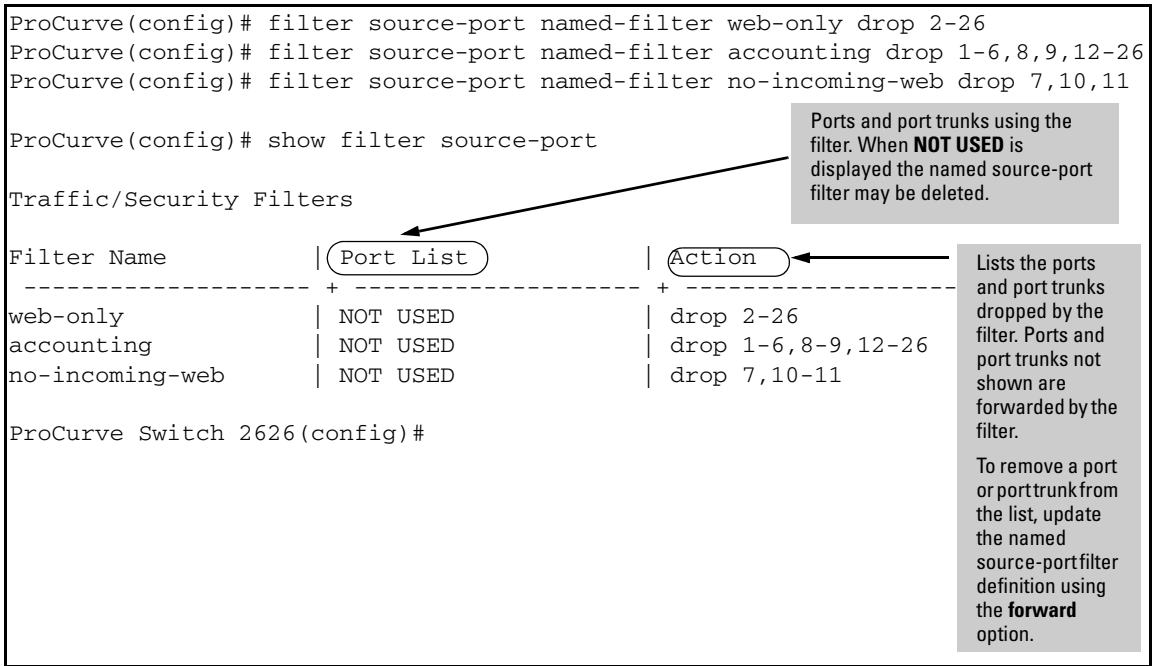


Figure 9-5. Applying Example Named Source-Port Filters

Once the named source-port filters have been defined and configured we now apply them to the switch ports.




Figure 9-6. Source Port Filters Applied to Switch Ports

The **show filter** command shows what ports have filters applied.


```
ProCurve(config)# show filter
```

Traffic/Security Filters

IDX	Filter Type	Value
1	Source Port	2
2	Source Port	3
3	Source Port	4
4	Source Port	5
5	Source Port	6
6	Source Port	8
7	Source Port	9
8	Source Port	12
20	Source Port	24
21	Source Port	25
22	Source Port	26
23	Source Port	7
24	Source Port	10
25	Source Port	11
26	Source Port	1

Indicates the port number or port-trunk name of the source port or trunk assigned to the filter.

An automatically assigned index number used to identify the filter for a detailed information listing. A filter retains its assigned IDX number for as long as the filter exists in the switch. The switch assigns the lowest available IDX number to a new filter. This can result in a newer filter having a lower IDX number than an older filter if a previous (source-port or named source-port) filter deletion created a gap in the filter listing.

Figure 9-7. Example of the show filter Command

Using the **IDX** value in the **show filter** command, we can see how traffic is filtered on a specific port (**Value**).The two outputs below show a non-accounting and an accounting switch port.

ProCurve(config)# show filter 4			ProCurve(config)# show filter 24			
Traffic/Security Filters			Traffic/Security Filters			
Filter Type : Source Port			Filter Type : Source Port			
Source Port : 5			Source Port : 10			
Dest	Port Type	Action	Dest	Port Type	Action	
-----		+	-----		+	-----
1	10/100TX	Forward	1	10/100TX	Drop	
2	10/100TX	Drop	2	10/100TX	Drop	
3	10/100TX	Drop	3	10/100TX	Drop	
4	10/100TX	Drop	4	10/100TX	Drop	
5	10/100TX	Drop	5	10/100TX	Drop	
6	10/100TX	Drop	6	10/100TX	Drop	
7	10/100TX	Drop	7	10/100TX	Forward	
8	10/100TX	Drop	8	10/100TX	Drop	
9	10/100TX	Drop	9	10/100TX	Drop	
10	10/100TX	Drop	10	10/100TX	Drop	
11	10/100TX	Drop	11	10/100TX	Drop	
12	10/100TX	Drop	12	10/100TX	Drop	
.	

Figure 9-8. Example Showing Traffic Filtered on Specific Ports

The same command, using IDX 26, shows how traffic from the Internet is handled.

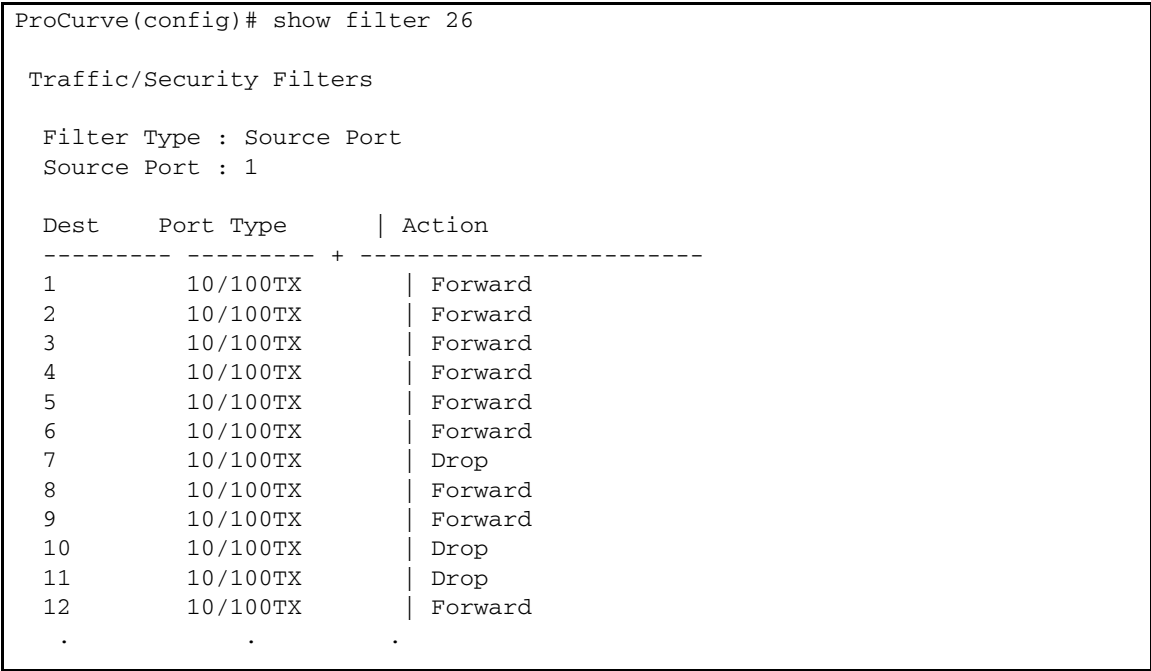


Figure 9-9. Example of Source Port Filtering with Internet Traffic

As the company grows, more resources are required in accounting. Two additional accounting workstations are added and attached to ports 12 and 13. A second server is added attached to port8.

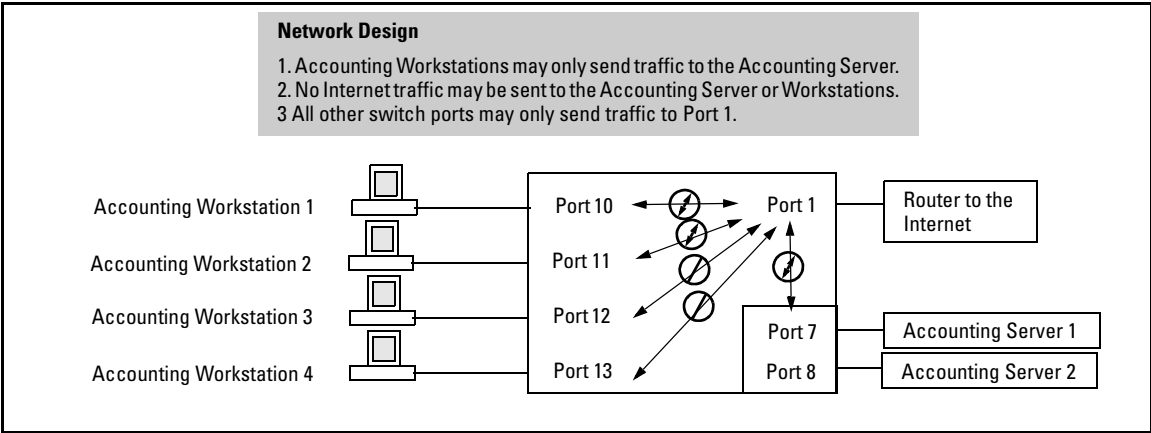


Figure 9-10. Expanded Network Configuration for Named Source-Port Filters Example

The following revisions to the named source-port filter definitions maintain the desired network traffic management, as shown in the **Action** column of the **show** command.

```
ProCurve(config)# filter source-port named-filter accounting forward 8,12,13
ProCurve(config)# filter source-port named-filter no-incoming-web drop 8,12,13
ProCurve(config)#
ProCurve(config)# show filter source-port
```

Traffic/Security Filters

Filter Name	Port List	Action
web-only	2-6,8-9,12-26	drop 2-26
accounting	7,10-11	drop 1-6,9,14-26
no-incoming-web	1	drop 7-8,10-13

```
ProCurve(config)#
```

Figure 9-11. Example Showing Network Traffic Management with Source Port Filters

We next apply the updated named source-port filters to the appropriate switch ports. As a port can only have one source-port filter (named or not named), before applying the new named source-port filters we first remove the existing source-port filters on the port.

```
ProCurve(config)# no filter source-port 8,12,13
ProCurve(config)# filter source-port 8,12,13 named-filter accounting
ProCurve(config)#
```

The named source-port filters now manage traffic on the switch ports as shown below, using the **show filter source-port** command.

```
ProCurve(config)# show filter source-port

Traffic/Security Filters

Filter Name          | Port List          | Action
-----+-----+-----
web-only             | 2-6,9,14-26        | drop 2-26
accounting           | 7-8,10-13          | drop 1-6,9,14-26
no-incoming-web      | 1                  | drop 7-8,10-13

ProCurve(config)#
```

Figure 9-12. Named Source-Port Filters Managing Traffic

Configuring Traffic/Security Filters

Use this procedure to specify the type of filters to use on the switch and whether to forward or drop filtered packets for each filter you specify.

1. Select the static filter type(s).
2. For inbound traffic matching the filter type, determine the filter action you want for each outbound (destination) port on the switch (forward or drop). The default action for a new filter is to forward traffic of the specified type to all outbound ports.
3. Configure the filter.
4. Use **show filter** (page 9-20) to check the filter listing to verify that you have configured correct action for the desired outbound ports.

Configuring a Source-Port Traffic Filter

Syntax: [no] filter

[source-port < port-number | trunk-name >]

*Specifies one inbound port or trunk. Traffic received inbound on this interface from other devices will be filtered. The **no** form of the command deletes the source-port filter for <port-number> and returns the destination ports for that filter to the **Forward** action. (Default: Forward on all ports.)*

Note: *If multiple VLANs are configured, the source-port and the destination port(s) must be in the same VLAN unless routing is enabled. Similarly, if a VLAN containing both the source and destination is multi-netted, the source and destination ports and/or trunks must be in the same subnet unless routing is enabled.*

[drop] < destination-port-list > [forward < port-list >]

*Configures the filter to drop traffic for the ports and/or trunks in the designated < destination-port-list >. Can be followed by **forward** < destination-port-list > if you have other destination ports set to **drop** that you want to change to **forward**. If no drop or forward action is specified, the switch automatically creates a filter with a **forward** action from the designated source port (or trunk) to all destination ports (or trunks) on the switch.*

[forward] < port-list >

*Configures the filter to forward traffic for the ports and/or trunks in the designated < destination-port-list >. Because **forward** is the default state for destinations in a filter, this command is useful when destinations in an existing filter are configured for **drop** and you want to change them to **forward**. Can be followed by **drop** < destination-port-list > if you have other destination ports set to **forward** that you want to change to **drop**. If no drop or forward action is specified, the switch automatically creates a filter with a forward action from the designated source port (or trunk) to all destination ports (or trunks) on the switch.*

Example of Creating a Source-Port Filter

For example, assume that you want to create a source-port filter that drops all traffic received on port 5 with a destination of port trunk 1 (**Trk1**) and any port in the range of port 10 to port 15. To create this filter you would execute this command:

```
ProCurve(config)# filter source-port 5 drop trk1,10-15
```

Later, suppose you wanted to shift the destination port range for this filter up by two ports; that is, to have the filter drop all traffic received on port 5 with a destination of any port in the range of port 12 to port 17. (The **Trk1** destination is already configured in the filter and can remain as-is.) With one command you can restore forwarding to ports 10 and 11 while adding ports 16 and 17 to the "drop" list:

```
ProCurve(config)# filter source-port 5 forward 10-11 drop  
16-17
```

Configuring a Filter on a Port Trunk

This operation uses the same command as is used for configuring a filter on an individual port. However, the configuration process requires two steps:

1. Configure the port trunk.
2. Configure a filter on the port trunk by using the trunk name (**trk1**, **trk2**, ...**trk6**) instead of a port name.

For example, to create a filter on port trunk 1 to drop traffic received inbound for trunk 2 and ports 10-15:

```
ProCurve(config)# filter source-port trk1 drop trk2,10-15
```

Note that if you first configure a filter on a port and then later add the port to a trunk, the port remains configured for filtering *but the filtering action will be suspended while the port is a member of the trunk*. That is, the trunk does not adopt filtering from the port configuration. You must still explicitly configure the filter on the port trunk. If you use the **show filter < index >** command for a filter created before the related source port was added to a trunk, the port number appears between asterisks (*), indicating that the filter action has been suspended for that filter. For example, if you create a

filter on port 5, then create a trunk with ports 5 and 6, and display the results, you would see the following:

```
ProCurve(config)# filter source-port 5 drop 2
ProCurve(config)# trunk 5-6 trkl
ProCurve(config)# show filter
```

Traffic/Security Filters

IDX	Filter Type	Value
1	Source Port	*5*

```
ProCurve(config)# show filter 1
```

Traffic/Security Filters

Filter Type : Source Port
Source Port : *5*

Dest Port	Type	Action
1	100/1000T	Forward
2	100/1000T	Drop
3	100/1000T	Forward
4	100/1000T	Forward
.	.	.
.	.	.
.	.	.

The *5* shows that port 5 is configured for filtering, but the filtering action has been suspended while the port is a member of a trunk.

If you want the trunk to which port 5 belongs to filter traffic, then you must explicitly configure filtering on the trunk.

Note: If you configure an existing trunk for filtering and later add another port to the trunk, the switch will apply the filter to all traffic moving on any link in the trunk. If you remove a port from the trunk it returns to the configuration it had before it was added to the trunk

Figure 9-13. Example of Switch Response to Adding a Filtered Source Port to a Trunk

Editing a Source-Port Filter

The switch includes in one filter the action(s) for all destination ports and/or trunks configured for a given source port or trunk. Thus, if a source-port filter already exists and you want to change the currently configured action for some destination ports or trunks, use the **filter source-port** command to update the existing filter. For example, suppose you configure a filter to drop traffic received on port 8 and destined for ports 1 and 2. The resulting filter is shown on the left in figure 9-14. Later, you update the filter to drop traffic received on port 8 and destined for ports 3 through 5. Since only one filter exists for a given source port, the filter on traffic from port 8 appears as shown on the right in figure 9-14:

ProCurve(config)# show filter 1				ProCurve(config)# show filter 1			
Traffic/Security Filters				Traffic/Security Filters			
Filter Type : Source Port				Filter Type : Source Port			
Source Port : 8				Source Port : 8			
Dest	Port	Type	Action	Dest	Port	Type	Action
1	100	/1000T	Drop	1	100	/1000T	Drop
2	100	/1000T	Drop	2	100	/1000T	Drop
3	100	/1000T	Forward	3	100	/1000T	Drop
4	100	/1000T	Forward	4	100	/1000T	Drop
5	100	/1000T	Forward	5	100	/1000T	Drop
6	100	/1000T	Forward	6	100	/1000T	Forward
7	100	/1000T	Forward	7	100	/1000T	Forward
8	100	/1000T	Forward	8	100	/1000T	Forward
9	100	/1000T	Forward	9	100	/1000T	Forward
10	100	/1000T	Forward	10	100	/1000T	Forward

Figure 9-14. Assigning Additional Destination Ports to an Existing Filter

For example, suppose you wanted to configure the filters in table 9-2 on a switch. (For more on source-port filters, refer to “Configuring a Source-Port Traffic Filter” on page 9-16.)

Table 9-2. Filter Example

Filter Type	Filter Value	Action	Destination Ports
Source-Port	Inbound ports: A1, A2*	Drop	D1-D4

*Because the switch allows one inbound port in a source-port filter, the requirement to filter ports A1 and A2 means you will configure two separate source-port filters.

The following commands configure the filters listed above:

```
ProCurve(config)# filter source-port a1 drop e d1-d4
ProCurve(config)# filter source-port a2 drop d1-d4
```

Figure 9-15. Configuring Various Traffic/Security Filters

Filter Indexing

The switch automatically assigns each new filter to the lowest-available index (IDX) number. The index numbers are included in the **show filter** command described in the next section and are used with the **show filter < index >** command to display detailed information about a specific filter.

If there are no filters currently configured, and you create three filters in succession, they will have index numbers 1 - 3. However, if you then delete the filter using index number “2” and then configure two new filters, the first

new filter will receive the index number "2" and the second new filter will receive the index number "4". This is because the index number "2" was made vacant by the earlier deletion, and was therefore the lowest index number available for the next new filter.

Displaying Traffic/Security Filters

This command displays a listing of all filters by index number and also enables you to use the index number to display the details of individual filters.

Syntax: show filter

Lists the filters configured in the switch, with corresponding filter index (IDX) numbers.

IDX: *An automatically assigned index number used to identify the filter for a detailed information listing. A filter retains its assigned IDX number for as long as the filter exists in the switch. The switch assigns the lowest available IDX number to a new filter. This can result in a newer filter having a lower IDX number than an older filter if a previous filter deletion created a gap in the filter listing.*

Filter Type: *Indicates the type of filter assigned to the IDX number (source-port).*

Value: *Indicates the port number or port-trunk name of the source port or trunk assigned to the filter*

[index]

Lists the filter type and other data for the filter corresponding to the index number in the **show filter** output. Also lists, for each outbound destination port in the switch, the port number, port type, and filter action (forward or drop). The switch assigns the lowest available index number to a new filter. If you delete a filter, the index number for that filter becomes available for the next filter you create.

Configuring Port-Based and User-Based Access Control (802.1X)

Contents

Overview	10-3
Why Use Port-Based or User-Based Access Control?	10-3
General Features	10-3
User Authentication Methods	10-4
802.1X User-Based Access Control	10-4
802.1X Port-Based Access Control	10-5
Alternative To Using a RADIUS Server	10-6
Accounting	10-6
Terminology	10-6
General 802.1X Authenticator Operation	10-9
Example of the Authentication Process	10-9
VLAN Membership Priority	10-10
General Operating Rules and Notes	10-12
General Setup Procedure for 802.1X Access Control	10-14
Do These Steps Before You Configure 802.1X Operation	10-14
Overview: Configuring 802.1X Authentication on the Switch	10-17
Configuring Switch Ports as 802.1X Authenticators	10-18
1. Enable 802.1X Authentication on Selected Ports	10-19
A. Enable the Selected Ports as Authenticators and Enable the (Default) Port-Based Authentication	10-19
B. Specify User-Based Authentication or Return to Port-Based Authentication	10-20
Example: Configuring User-Based 802.1X Authentication	10-21
Example: Configuring Port-Based 802.1X Authentication	10-21
2. Reconfigure Settings for Port-Access	10-21

3. Configure the 802.1X Authentication Method	10-24
4. Enter the RADIUS Host IP Address(es)	10-25
5. Enable 802.1X Authentication on the Switch	10-25
6. Optional: Reset Authenticator Operation	10-26
7. Optional: Configure 802.1X Controlled Directions	10-26
Wake-on-LAN Traffic	10-27
Operating Notes	10-27
Example: Configuring 802.1X Controlled Directions	10-28
802.1X Open VLAN Mode	10-29
Introduction	10-29
VLAN Membership Priorities	10-30
Use Models for 802.1X Open VLAN Modes	10-31
Operating Rules for Authorized-Client and Unauthorized-Client VLANs	10-36
Setting Up and Configuring 802.1X Open VLAN Mode	10-40
802.1X Open VLAN Operating Notes	10-44
Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1X-Authenticated Devices	10-45
Port-Security	10-46
Configuring Switch Ports To Operate As Supplicants for 802.1X Connections to Other Switches	10-47
Example	10-47
Supplicant Port Configuration	10-49
Displaying 802.1X Configuration, Statistics, and Counters	10-51
Show Commands for Port-Access Authenticator	10-51
Viewing 802.1X Open VLAN Mode Status	10-61
Show Commands for Port-Access Supplicant	10-65
How RADIUS/802.1X Authentication Affects VLAN Operation .	10-66
VLAN Assignment on a Port	10-67
Operating Notes	10-67
Example of Untagged VLAN Assignment in a RADIUS-Based Authentication Session	10-69
Enabling the Use of GVRP-Learned Dynamic VLANs in Authentication Sessions	10-72
Messages Related to 802.1X Operation	10-74

Overview

Feature	Default	Menu	CLI	Web
Configuring Switch Ports as 802.1X Authenticators	Disabled	n/a	page 10-18	n/a
Configuring 802.1X Open VLAN Mode	Disabled	n/a	page 10-29	n/a
Configuring Switch Ports to Operate as 802.1X Supplicants	Disabled	n/a	page 10-47	n/a
Displaying 802.1X Configuration, Statistics, and Counters	n/a	n/a	page 10-51	n/a
How 802.1X Affects VLAN Operation	n/a	n/a	page 10-65	n/a
RADIUS Authentication and Accounting	Refer to chapter 5, "RADIUS Authentication, Authorization, and Accounting"			

Why Use Port-Based or User-Based Access Control?

Local Area Networks are often deployed in a way that allows unauthorized clients to attach to network devices, or allows unauthorized users to get access to unattended clients on a network. Also, the use of DHCP services and zero configuration make access to networking services easily available. This exposes the network to unauthorized use and malicious attacks. While access to the network should be made easy, uncontrolled and unauthorized access is usually not desirable. 802.1X simplifies security management by providing access control along with the ability to control user profiles from up to three RADIUS servers while allowing a given user to use the same entering valid user credentials for access from multiple points within the network.

General Features

802.1X on the switches covered in this guide includes the following:

- Switch operation as both an authenticator (for supplicants having a point-to-point connection to the switch) and as a supplicant for point-to-point connections to other 802.1X-aware switches.
 - Authentication of 802.1X access using a RADIUS server and either the EAP or CHAP protocol.
 - Provision for enabling clients that do not have 802.1 supplicant software to use the switch as a path for downloading the software and initiating the authentication process (802.1X Open VLAN mode).
 - User-Based access control option with support for up to 32 authenticated clients per-port.

- Port-Based access control option allowing authentication by a single client to open the port. This option does not force a client limit and, on a port opened by an authenticated client, allows unlimited client access without requiring further authentication.
- Supplicant implementation using CHAP authentication and independent user credentials on each port.
- The local operator password configured with the **password** command for management access to the switch is no longer accepted as an 802.1X authenticator credential. The **password port-access** command configures the local operator username and password used as 802.1X authentication credentials for access to the switch. The values configured can be stored in a configuration file using the **include-credentials** command. For information about the **password port-access** command, see “Do These Steps Before You Configure 802.1X Operation” on page 10-14.
- On-demand change of a port’s configured VLAN membership status to support the current client session.
- Session accounting with a RADIUS server, including the accounting update interval.
- Use of Show commands to display session counters.
- Support for concurrent use of 802.1X and either Web authentication or MAC authentication on the same port.
- For unauthenticated clients that do not have the necessary 802.1X supplicant software (or for other reasons related to unauthenticated clients), there is the option to configure an *Unauthorized-Client VLAN*. This mode allows you to assign unauthenticated clients to an isolated VLAN through which you can provide the necessary supplicant software and/or other services you want to extend to these clients.

User Authentication Methods

The switch offers two methods for using 802.1X access control. Generally, the “Port Based” method supports one 802.1X-authenticated client on a port, which opens the port to an unlimited number of clients. The “User-Based” method supports up to 32 802.1X-authenticated clients on a port. In both cases, there are operating details to be aware of that can influence your choice of methods.

802.1X User-Based Access Control

802.1X operation with access control on a per-user basis provides client-level security that allows LAN access to individual 802.1X clients (up to 32 per port), where each client gains access to the LAN by entering valid user credentials.

This operation improves security by opening a given port only to individually authenticated clients, while simultaneously blocking access to the same port for clients that cannot be authenticated. All sessions must use the same untagged VLAN. Also, an authenticated client can use any tagged VLAN memberships statically configured on the port, provided the client is configured to use the tagged VLAN memberships available on the port. (Note that the session total includes any sessions begun by the Web Authentication or MAC Authentication features covered in chapter 3.) For more information, refer to “Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1X-Authenticated Devices” on page 10-45.

802.1X Port-Based Access Control

802.1X port-based access control provides port-level security that allows LAN access only on ports where a single 802.1X-capable client (supplicant) has entered authorized RADIUS user credentials. For reasons outlined below, this option is recommended for applications where only one client at a time can connect to the port. Using this option, the port processes all traffic as if it comes from the same client. Thus, in a topology where multiple clients can connect to the same port at the same time:

- If the first client authenticates and opens the port, and then another client authenticates, the port responds as if the original client has initiated a reauthentication. With multiple clients authenticating on the port, the RADIUS configuration response to the latest client authentication replaces any other configuration from an earlier client authentication. If all clients use the same configuration this should not be a problem. But if the RADIUS server responds with different configurations for different clients, then the last client authenticated will effectively lock out any previously authenticated client. When *any* client to authenticate closes its session, the port will also close and remain so until another client successfully authenticates.
- The most recent client authentication determines the untagged VLAN membership for the port. Also, any client able to use the port can access any tagged VLAN memberships statically configured on the port, provided the client is configured to use the available, tagged VLAN memberships.
- If the first client authenticates and opens the port, and then one or more other clients connect without trying to authenticate, then the port configuration as determined by the original RADIUS response remains unchanged and all such clients will have the same access as the authenticated client. When the authenticated client closes the session, the port will also be closed to any other, unauthenticated clients that may have also been using the port.

This operation unblocks the port while an authenticated client session is in progress. In topologies where simultaneous, multiple client access is possible this can allow unauthorized and unauthenticated access by another client while an authenticated client is using the port. If you want to allow only authenticated clients on the port, then user-based access control (page 10-4) should be used instead of port-based access control. Using the user-based method enables you to specify up to 32 authenticated clients.

Note

Port-Based 802.1X can operate concurrently with Web-Authentication or MAC-Authentication on the same port. However, this is not a commonly used application and is not generally recommended. For more information, refer to the operating note on page 10-13.

Alternative To Using a RADIUS Server

Note that you can also configure 802.1X for authentication through the switch's local username and password instead of a RADIUS server, but doing so increases the administrative burden, decentralizes user credential administration, and reduces security by limiting authentication to one Operator password set for all users.

Accounting

The switches covered in this guide also provide RADIUS Network accounting for 802.1X access. Refer to chapter 5, "RADIUS-Administered CoS and Rate-Limiting".

Terminology

802.1X-Aware: Refers to a device that is running either 802.1X authenticator software or 802.1X client software and is capable of interacting with other devices on the basis of the IEEE 802.1X standard.

Authorized-Client VLAN: Like the Unauthorized-Client VLAN, this is a conventional, static VLAN previously configured on the switch by the System Administrator. The intent in using this VLAN is to provide authenticated clients with network services that are not available on either the port's statically configured VLAN memberships or any VLAN memberships that may be assigned during the RADIUS authentication process. While an 802.1X port is a member of this VLAN, the port is untagged. When

a port loses its authenticated client connection, it drops its membership in this VLAN. Note that with multiple clients on a port, all such clients use the same untagged, port-based VLAN membership.

Authentication Server: The entity providing an authentication service to the switch when the switch is configured to operate as an authenticator. In the case of a switch running 802.1X, this is a RADIUS server (unless local authentication is used, in which case the switch performs this function using its own username and password for authenticating a supplicant).

Authenticator: In ProCurve applications, a switch that requires a supplicant to provide the proper credentials before being allowed access to the network.

CHAP (MD5): Challenge Handshake Authentication Protocol.

Client: In this application, an end-node device such as a management station, workstation, or mobile PC linked to the switch through a point-to-point LAN link.

User-Based Authentication: The 802.1X extension in the switches covered in this guide. In this operation, multiple clients on the same port must individually authenticate themselves.

Guest VLAN: See “Unauthorized-Client VLAN”.

EAP (Extensible Authentication Protocol): EAP enables network access that supports multiple authentication methods.

EAPOL: Extensible Authentication Protocol Over LAN, as defined in the 802.1X standard.

Friendly Client: A client that does not pose a security risk if given access to the switch and your network.

MD5: An algorithm for calculating a unique digital signature over a stream of bytes. It is used by CHAP to perform authentication without revealing the shared secret (password).

PVID (Port VID): This is the VLAN ID for the untagged VLAN to which an 802.1X port belongs.

Port-Based Authentication: In this operation, the first client on a port to authenticate itself unblocks the port for the duration of the client’s 802.1X-authenticated session. The switches covered in this guide use port-based authentication.

Static VLAN: A VLAN that has been configured as “permanent” on the switch by using the CLI **vlan < vid >** command or the Menu interface.

Supplicant: The entity that must provide the proper credentials to the switch before receiving access to the network. This is usually an end-user workstation, but it can be a switch, router, or another device seeking network services.

Tagged Membership in a VLAN: This type of VLAN membership allows a port to be a member of multiple VLANs simultaneously. If a client connected to the port has an operating system that supports 802.1Q VLAN tagging, then the client can access VLANs for which the port is a tagged member. If the client does not support VLAN tagging, then it can access only a VLAN for which the port is an untagged member. (A port can be an untagged member of only one port-based VLAN at a time.) Where a port is a tagged member of a VLAN, 802.1X Open VLAN mode does not affect the port’s access to the VLAN unless the port is statically configured as a member of a VLAN that is also configured as the Unauthorized-Client or Authorized-Client VLAN. See also “**Untagged Membership in a VLAN**”.

Unauthorized-Client VLAN: A conventional, static VLAN statically configured on the switch. It is used to provide access to a client prior to authentication, and is sometimes termed a *guest* VLAN. It should be set up to allow an unauthenticated client to access only the initialization services necessary to establish an authenticated connection, plus any other desirable services whose use by an unauthenticated client poses no security threat to your network. (Note that an unauthenticated client has access to all network resources that have membership in the VLAN you designate as the Unauthorized-Client VLAN.) A port configured to use a given Unauthorized-Client VLAN does not have to be statically configured as a member of that VLAN as long as at least one other port on the switch is statically configured as a tagged or untagged member of the same Unauthorized-Client VLAN. An unauthorized-client VLAN is available on a port only if there is no authenticated client already using the port.

Untagged Membership in a VLAN: A port can be an untagged member of only one VLAN. (In the factory-default configuration, all ports on the switch are untagged members of the default VLAN.) An untagged VLAN membership is *required* for a client that does not support 802.1Q VLAN tagging. A port can simultaneously have one untagged VLAN membership and multiple tagged VLAN memberships. Depending on how you configure 802.1X Open VLAN mode for a port, a statically configured, untagged VLAN membership may become unavailable while there is a client session on the port. See also “**Tagged Membership in a VLAN**”.

General 802.1X Authenticator Operation

This operation provides security on a point-to-point link between a client and the switch, where both devices are 802.1X-aware. (If you expect desirable clients that do not have the necessary 802.1X supplicant software, you can provide a path for downloading such software by using the 802.1X Open VLAN mode—refer to “802.1X Open VLAN Mode” on page 10-29.)

Example of the Authentication Process

Suppose that you have configured a port on the switch for 802.1X authentication operation, which blocks access to the LAN through that port. If you then connect an 802.1X-aware client (supplicant) to the port and attempt to log on:

1. The switch responds with an identity request.
2. The client responds with a user name that uniquely defines this request for the client.
3. The switch responds in one of the following ways:
 - If 802.1X on the switch is configured for RADIUS authentication, the switch then forwards the request to a RADIUS server.
 - i. The server responds with an access challenge which the switch forwards to the client.
 - ii. The client then provides identifying credentials (such as a user certificate), which the switch forwards to the RADIUS server.
 - iii. The RADIUS server then checks the credentials provided by the client.
 - iv. If the client is successfully authenticated and authorized to connect to the network, then the server notifies the switch to allow access to the client. Otherwise, access is denied and the port remains blocked.
 - If 802.1X on the switch is configured for local authentication, then:
 - i. The switch compares the client's credentials to the username and password configured in the switch (Operator level).
 - ii. If the client is successfully authenticated and authorized to connect to the network, then the switch allows access to the client. Otherwise, access is denied and the port remains blocked for that client.

Note

The switches covered in this guide can use either 802.1X port-based authentication or 802.1X user-based authentication. For more information, refer to “User Authentication Methods” on page 10-4.

VLAN Membership Priority

Following client authentication, an 802.1X port resumes membership in any tagged VLANs for which it is already assigned in the switch configuration. The port also becomes an untagged member of one VLAN according to the following order of options:

- a. **1st Priority:** The port joins a VLAN to which it has been assigned by a RADIUS server during client authentication.
- b. **2nd Priority:** If RADIUS authentication does not include assigning the port to a VLAN, then the switch assigns the port to the VLAN entered in the port’s 802.1X configuration as an *Authorized-Client* VLAN, if configured.
- c. **3rd Priority:** If the port does not have an Authorized-Client VLAN configured, but does have a static, untagged VLAN membership in its configuration, then the switch assigns the port to this VLAN.

A port assigned to a VLAN by an Authorized-Client VLAN configuration (or a RADIUS server) will be an untagged member of the VLAN for the duration of the authenticated session. This applies even if the port is also configured in the switch as a tagged member of the same VLAN.

Note

On the switches covered in this guide, using the same port for both RADIUS-assigned clients and clients using a configured, Authorized-Client VLAN is not recommended. This is because doing so can result in authenticated clients with mutually exclusive VLAN priorities, which means that some authenticated clients can be denied access to the port. Refer to figure 10-1 on page 10-11.

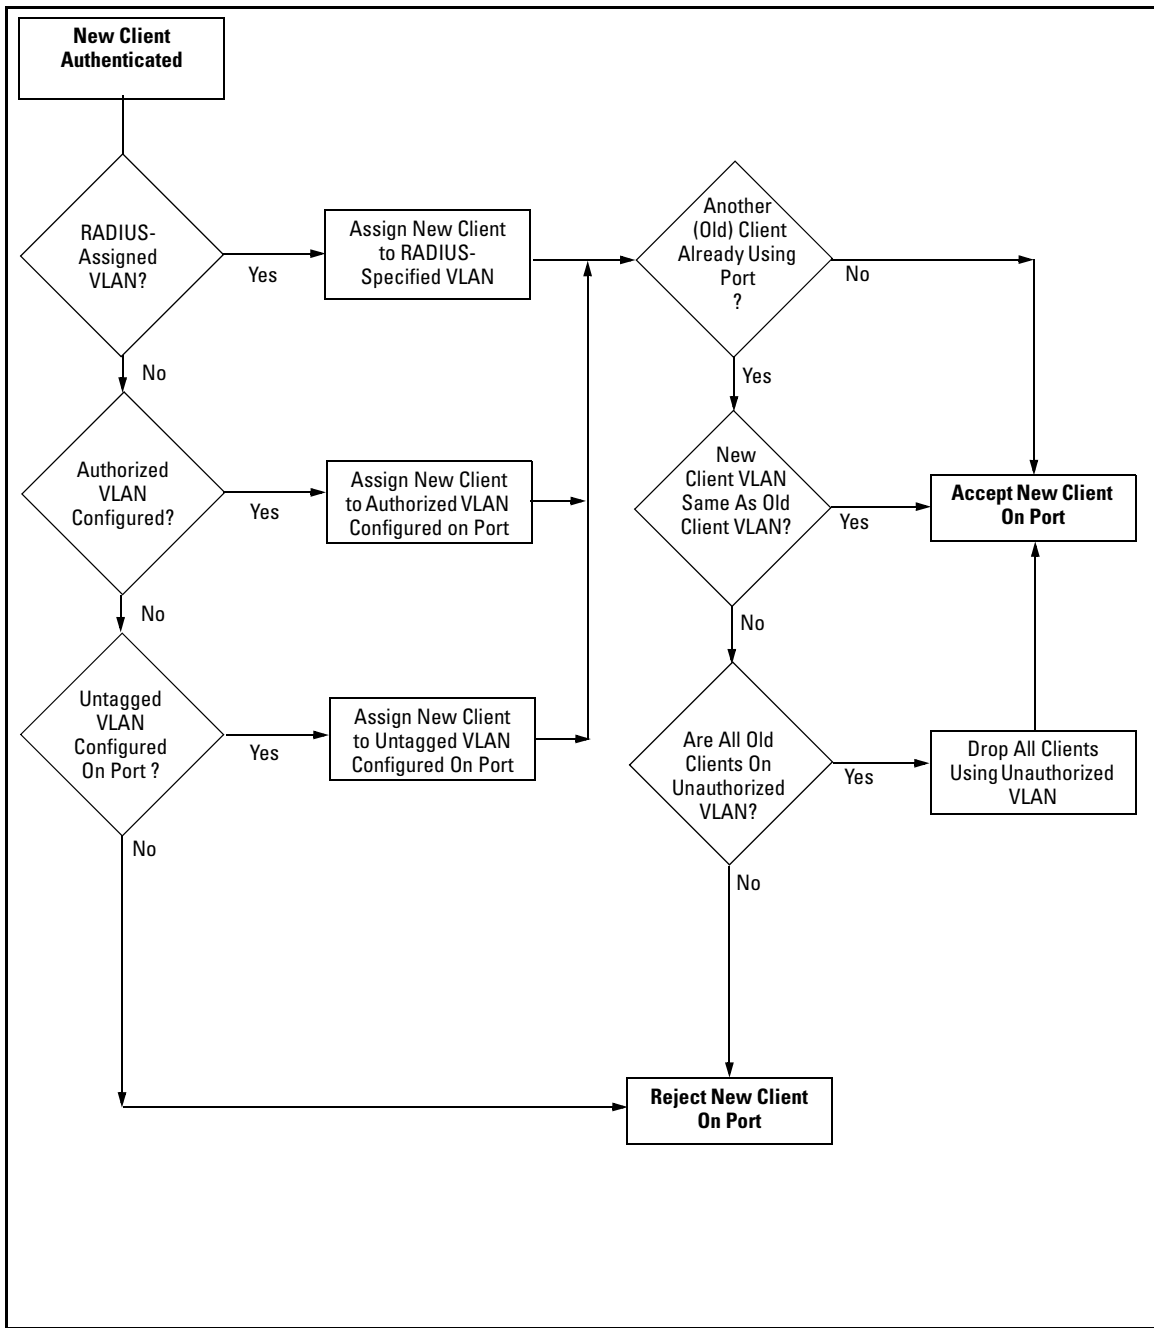


Figure 10-1. Priority of VLAN Assignment for an Authenticated Client

General Operating Rules and Notes

- In the user-based mode, when there is an authenticated client on a port, the following traffic movement is allowed:
 - Multicast and broadcast traffic is allowed on the port.
 - Unicast traffic to authenticated clients on the port is allowed.
 - All traffic from authenticated clients on the port is allowed.
- When a port on the switch is configured as either an authenticator or supplicant and is connected to another device, rebooting the switch causes a re-authentication of the link.
- Using user-based 802.1X authentication, when a port on the switch is configured as an authenticator the port allows only authenticated clients up to the currently configured client limit.

For clients that do not have the proper 802.1X supplicant software, the optional 802.1X Open VLAN mode can be used to open a path for downloading 802.1X supplicant software to a client or to provide other services for unauthenticated clients. Refer to “802.1X Open VLAN Mode” on page 10-29.)

- Using port-based 802.1X authentication, When a port on the switch is configured as an authenticator, one authenticated client opens the port. Other clients that are not running an 802.1X supplicant application can have access to the switch and network through the opened port. If another client uses an 802.1X supplicant application to access the opened port, then a re-authentication occurs using the RADIUS configuration response for the latest client to authenticate. To control access by all clients, use the user-based method.
- Where a switch port is configured with user-based authentication to accept multiple 802.1X (and/or Web- or MAC-Authentication) client sessions, all authenticated clients must use the same port-based, untagged VLAN membership assigned for the earliest, currently active client session. Thus, on a port where one or more authenticated client sessions are already running, all such clients will be on the same untagged VLAN. If a RADIUS server subsequently authenticates a new client, but attempts to re-assign the port to a different, untagged VLAN than the one already in use for the previously existing, authenticated client sessions, the connection for the new client will fail. For more on this topic, refer to “802.1X Open VLAN Mode” on page 10-29. (Note that if the port is statically configured with any tagged VLAN memberships, any authenticated client configured to use these tagged VLANs will have access to them.)

- If a port on switch “A” is configured as an 802.1X supplicant and is connected to a port on another switch, “B”, that is not 802.1X-aware, access to switch “B” will occur without 802.1X security protection.
- On a port configured for 802.1X with RADIUS authentication, if the RADIUS server specifies a VLAN for the supplicant and the port is a trunk member, the port will be blocked. If the port is later removed from the trunk, the port will allow authentication of the supplicant. Similarly, if the supplicant is authenticated and later the port becomes a trunk member, the port will be blocked. If the port is then removed from the trunk, it will allow the supplicant to re-authenticate.
- If a client already has access to a switch port when you configure the port for 802.1X authenticator operation, the port will block the client from further network access until it can be authenticated.
- Meshing is not supported on ports configured for 802.1X port-access security.
- A port can be configured as an authenticator *or* an 802.1X supplicant, or both. Some configuration instances block traffic flow or allow traffic to flow without authentication. Refer to “Configuring Switch Ports To Operate As Supplicants for 802.1X Connections to Other Switches” on page 10-47.
- To help maintain security, 802.1X and LACP cannot both be enabled on the same port. If you try to configure 802.1X on a port already configured for LACP (or the reverse) you will see a message similar to the following:

Error configuring port X: LACP and 802.1X cannot be run together.

Applying Web Authentication or MAC Authentication Concurrently with Port-Based 802.1X Authentication: While 802.1X port-based access control can operate concurrently with Web Authentication or MAC Authentication, port-based access control is subordinate to Web-Auth and MAC-Auth operation. If 802.1X operates in port-based mode and MAC or Web authentication is enabled on the same port, any 802.1X authentication has no effect on the ability of a client to access the controlled port. That is, the client’s access will be denied until the client authenticates through Web-Auth or MAC-Auth on the port. Note also that a client authenticating with port-based 802.1X does not open the port in the same way that it would if Web-Auth or MAC-Auth were not enabled. That is, any non-authenticating client attempting to access the port after another client authenticates with port-based 802.1X would still have to authenticate through Web-Auth or MAC-Auth.

General Setup Procedure for 802.1X Access Control

Do These Steps Before You Configure 802.1X Operation

1. Configure a local username and password on the switch for both the Operator (login) and Manager (enable) access levels. (While this may or may not be required for your 802.1X configuration, ProCurve recommends that you use a local username and password pair at least until your other security measures are in place.)

For switches covered in this guide, the local operator password configured with the password command is not accepted as an 802.1X authenticator credential. The port-access command is used to configure the operator username and password that are used as 802.1X credentials for network access to the switch. 802.1X network access is not allowed unless a password has been configured using the **password port-access** command.

Syntax: password port-access [user-name <name>] <password>

Configures the operator username and password used to access the network through 802.1X authentication.

user-name <name>

*Operator username (text string) used only for local authentication of 802.1X clients. This value is different from the local operator username configured with the **password** command for management access.*

<password>

*Operator password (text string) used only for local authentication of 802.1X clients. This value is different from the local operator password configured with the **password** command for management access.*

Figure 10-2 shows how to configure a local operator password for 802.1X access.


```
ProCurve(config)# password port-access user-name Jim secret3
```

Figure 10-2. Example of the Password Port-Access Command

You can save the port-access password for 802.1X authentication in the configuration file by using the **include-credentials** command. For more information, see “Saving Security Credentials in a Config File” on page 2-10. in this guide.

2. Determine the switch ports that you want to configure as authenticators and/or supplicants, and disable LACP on these ports. (For more information on disabling LACP, refer to the “Note” on page 10-19.).

To display the current configuration of 802.1X, Web-based, and MAC authentication on all switch ports, enter the **show port-access config** command.

```
ProCurve (config)# show port-access config

Port Access Status Summary
Port-access authenticator activated [No] : Yes
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : Yes

      Supplicant Authenticator Web Auth Mac Auth
Port Enabled   Enabled   Enabled   Enabled
-----
1      Yes      No        No        Yes
2      No       Yes       No        Yes
3      No       Yes       No        No
4      No       No        No        No
5      No       No        No        No
6      No       No        No        No
7      No       No        No        No
8      No       No        No        No
9      No       No        No        No
10     No       No        No        No
11     No       No        No        No
12     No       No        No        No
...

```

Figure 10-3. Example of show port-access config Command Output

Configuring Port-Based and User-Based Access Control (802.1X)

General Setup Procedure for 802.1X Access Control

3. Determine whether to use user-based access control (page 10-4) or port-based access control (page 10-5).
4. Determine whether to use the optional 802.1X Open VLAN mode for clients that are not 802.1X-aware; that is, for clients that are not running 802.1X supplicant software. (This will require you to provide downloadable software that the client can use to enable an authentication session.) For more on this topic, refer to “802.1X Open VLAN Mode” on page 10-29.
5. For any port you want to operate as a supplicant, determine the user credentials. You can either use the same credentials for each port or use unique credentials for individual ports or subgroups of ports. (This can also be the same local username/password pair that you assign to the switch.)
6. Unless you are using only the switch's local username and password for 802.1X authentication, configure at least one RADIUS server to authenticate access requests coming through the ports on the switch from external supplicants (including switch ports operating as 802.1X supplicants). You can use up to three RADIUS servers for authentication; one primary and two backups. Refer to the documentation provided with your RADIUS application.

Overview: Configuring 802.1X Authentication on the Switch

This section outlines the steps for configuring 802.1X on the switch. For detailed information on each step, refer to the following:

- “802.1X User-Based Access Control” on page 10-4
 - “802.1X Port-Based Access Control” on page 10-5
 - “Configuring Switch Ports To Operate As Supplicants for 802.1X Connections to Other Switches” on page 10-47.
1. Enable 802.1X user-based or port-based authentication on the individual ports you want to serve as authenticators. On the ports you will use as authenticators, either accept the default 802.1X settings or change them, as necessary. Note that, by default, the port-control parameter is set to **auto** for all ports on the switch. This requires a client to support 802.1X authentication and to provide valid credentials to get network access. Refer to page 10-19.
 2. If you want to provide a path for clients without 802.1X supplicant software to download the software so that they can initiate an authentication session, enable the 802.1X Open VLAN mode on the ports you want to support this feature. Refer to page 10-29.
 3. Configure the 802.1X authentication type. Options include:
 - Local Operator username and password (using the **password port-access** command).
 - EAP RADIUS: This option requires your RADIUS server application to support EAP authentication for 802.1X.
 - CHAP (MD5) RADIUS: This option requires your RADIUS server application to support CHAP (MD5) authentication.Refer to page 10-24.
 4. If you select either **eap-radius** or **chap-radius** for step 3, use the **radius host** command to configure up to three RADIUS server IP address(es) on the switch. See page 10-25.
 5. Enable 802.1X authentication on the switch. Refer to “1. Enable 802.1X Authentication on Selected Ports” on page 10-19.
 6. Test both the authorized and unauthorized access to your system to ensure that the 802.1X authentication works properly on the ports you have configured for port-access.

Note

If you want to implement the optional port security feature (step 7) on the switch, you should first ensure that the ports you have configured as 802.1X authenticators operate as expected.

7. If you are using Port Security on the switch, configure the switch to allow only 802.1X access on ports configured for 802.1X operation, and (if desired) the action to take if an unauthorized device attempts access through an 802.1X port. Refer to page 10-45.
8. If you want a port on the switch to operate as a supplicant on a port operating as an 802.1X authenticator on another device, then configure the supplicant operation. (Refer to “Configuring Switch Ports To Operate As Supplicants for 802.1X Connections to Other Switches” on page 10-47.)

Configuring Switch Ports as 802.1X Authenticators

802.1X Authentication Commands	Page
[no] aaa port-access authenticator < <i>port-list</i> >	10-19
[auth-vid clear-statistics client-limit control max-requests initialize logoff-period quiet-period server-timeout reauthenticate reauth-period supplicant-timeout tx-period unauth-period unauth-vid]	10-19
aaa authentication port-access < local eap-radius chap-radius >	10-24
[no] aaa port-access authenticator active	10-18
aaa port-access < <i>port-list</i> > controlled-directions < both in >	10-26
[no] port-security [ethernet] < <i>port-list</i> > learn-mode port-access	10-45
802.1X Open VLAN Mode Commands	10-29
802.1X Supplicant Commands	10-47
802.1X-Related Show Commands	10-51
RADIUS server configuration	10-25

1. Enable 802.1X Authentication on Selected Ports

This task configures the individual ports you want to operate as 802.1X authenticators for point-to-point links to 802.1X-aware clients or switches, and consists of two steps:

- A. Enable the selected ports as authenticators.
- B. Specify either user-based or port-based 802.1X authentication.

(Actual 802.1X operation does not commence until you perform step 5 on page 10-25 to activate 802.1X authentication on the switch.)

Note

If you enable 802.1X authentication on a port, the switch automatically disables LACP on that port. However, if the port is already operating in an LACP trunk, you must remove the port from the trunk before you can configure it for 802.1X authentication.

A. Enable the Selected Ports as Authenticators and Enable the (Default) Port-Based Authentication

Syntax: [no] aaa port-access authenticator < port-list >

*Enables specified ports to operate as 802.1X authenticators and enables port-based authentication. (To enable user-based authentication, execute this command first, and then execute the client-limit < port-list > version of this command described in the next section.) The **no** form of the command removes 802.1X authentication from < port-list >. To activate configured 802.1X operation, you must enable 802.1X authentication. Refer to “5. Enable 802.1X Authentication on the switch” on page 10-25.*

B. Specify User-Based Authentication or Return to Port-Based Authentication

User-Based 802.1X Authentication.

Syntax: `aaa port-access authenticator client-limit < port-list > < 1 - 32 >`

*Used after executing **aaa port-access authenticator < port-list >** (above) to convert authentication from port-based to user-based. Specifies user-based 802.1X authentication and the maximum number of 802.1X-authenticated client sessions allowed on each of the ports in **< port-list >**. If a port currently has no authenticated client sessions, the next authenticated client session the port accepts determines the untagged VLAN membership to which the port is assigned during the session. If another client session begins later on the same port while an earlier session is active, the later session will be on the same untagged VLAN membership as the earlier session.*

Note: *Because a switch allows 802.1X authentication and Web or MAC authentication to co-exist on the same port, the sum of authenticated client sessions allowed on a given port for both 802.1X and either Web- or MAC-authentication cannot exceed 32.*

Port-Based 802.1X Authentication.

`no aaa port-access authenticator client-limit`

*Used to convert a port from user-based authentication to port-based authentication, which is the default setting for ports on which authentication is enabled. (Executing **aaa port-access authenticator < port-list >** enables 802.1X authentication on **< port-list >** and enables port-based authentication—page 10-19.) If a port currently has no authenticated client sessions, the next authenticated client session the port accepts determines the untagged VLAN membership to which the port is assigned during the session. If another authenticated client session begins later on the same port while an earlier session is active, the later session replaces the currently active session and will be on the untagged VLAN membership specified by the RADIUS server for the later session.*

Example: Configuring User-Based 802.1X Authentication

This example enables ports A10-A12 to operate as authenticators, and then configures the ports for user-based authentication.

```
ProCurve(config)# aaa port-access authenticator a10-A12
ProCurve(config)# aaa port-access authenticator a10-A12 client-limit 4
```

Figure 10-4. Example of Configuring User-Based 802.1X Authentication

Example: Configuring Port-Based 802.1X Authentication

This example enables ports A13-A15 to operate as authenticators, and then configures the ports for port-based authentication.

```
ProCurve(config)# aaa port-access authenticator a13-a15
ProCurve(config)# no aaa port-access authenticator a13-a15 client-limit
```

Figure 10-5. Example of Configuring Port-Based 802.1X Authentication

2. Reconfigure Settings for Port-Access

The commands in this section are initially set by default and can be reconfigured as needed.

Syntax: aaa port-access authenticator < *port-list* >
[control < authorized | auto | unauthorized >]

Controls authentication mode on the specified port:

authorized: Also termed “*Force Authorized*”. Gives access to a device connected to the port. In this case, the device does not have to provide 802.1X credentials or support 802.1X authentication. (You can still configure console, Telnet, or SSH security on the port.)

auto (the default): The device connected to the port must support 802.1X authentication and provide valid credentials to get network access. (Optional: You can use the Open VLAN mode to provide a path for clients without 802.1X supplicant software to down-load this software and begin the authentication process. Refer to “802.1X Open VLAN Mode” on page 10-29.)

unauthorized: Also termed "Force Unauthorized". Do not grant access to the network, regardless of whether the device provides the correct credentials and has 802.1X support. In this state, the port blocks access to any connected device.

[quiet-period < 0 - 65535 >]

*Sets the period during which the port does not try to acquire a supplicant. The period begins after the last attempt authorized by the **max-requests** parameter fails (next page). (Default: 60 seconds)*

[tx-period < 0 - 65535 >]

Sets the period the port waits to retransmit the next EAPOL PDU during an authentication session. (Default: 30 seconds)

— Continued —

aaa port-access authenticator < port-list >

[supplicant-timeout < 1 - 300 >]

Sets the period of time the switch waits for a supplicant response to an EAP request. If the supplicant does not respond within the configured time frame, the session times out. (Default: 30 seconds)

[server-timeout < 1 - 300 >]

*Sets the period of time the switch waits for a server response to an authentication request. If there is no response within the configured time frame, the switch assumes that the authentication attempt has timed out. Depending on the current **max-requests** setting, the switch will either send a new request to the server or end the authentication session. (Default: 30 seconds)*

[max-requests < 1 - 10 >]

*Sets the number of authentication attempts that must time-out before authentication fails and the authentication session ends. If you are using the Local authentication option, or are using RADIUS authentication with only one host server, the switch will not start another session until a client tries a new access attempt. If you are using RADIUS authentication with two or three host servers, the switch will open a session with each server, in turn, until authentication occurs or there are no more servers to try. During the **quiet-period** (previous page), if any, you cannot reconfigure this parameter. (Default: 2)*

—Continued—

[reauth-period < 0 - 9999999 >]

Sets the period of time after which clients connected must be re-authenticated. When the timeout is set to 0 the reauthentication is disabled (Default: 0 second)

[unauth-vid < vlan-id >]

Configures an existing static VLAN to be the Unauthorized-Client VLAN. This enables you to provide a path for clients without supplicant software to download the software and begin an authentication session. Refer to “802.1X Open VLAN Mode” on page 10-29.

aaa port-access authenticator < port-list >

[logoff-period]< 1 - 999999999 >

Configures the period of time the switch waits for client activity before removing an inactive client from the port. (Default: 300 seconds)

[unauth-period < 0-255 >]

Specifies a delay in seconds for placing a port on the Unauthorized-Client VLAN. This delay allows more time for a client with 802.1X supplicant capability to initiate an authentication session. If a connected client does not initiate a session before the timer expires, the port is assigned to the Unauthenticated-Client VLAN. (Default: 0 seconds)

[auth-vid < vid >]

Configures an existing, static VLAN to be the Authorized-Client VLAN. Refer to “802.1X Open VLAN Mode” on page 10-29.

3. Configure the 802.1X Authentication Method

This task specifies how the switch authenticates the credentials provided by a supplicant connected to a switch port configured as an 802.1X authenticator.

You can configure **local**, **chap-radius** or **eap-radius** as the primary password authentication method for the port-access method. You also need to select **none** or **authorized** as a secondary, or backup, method.

Syntax: `aaa authentication port-access <chap-radius leap-radius | local>`

*Configures **local**, **chap-radius** (MD5), or **eap-radius** as the primary password authentication method for port-access. The default primary authentication is **local**. (Refer to the documentation for your RADIUS server application.)*

*For switches covered in this guide, you must use the **password port-access** command to configure the operator username and password for 802.1X access. See “General Setup Procedure for 802.1X Access Control” on page 10-14 for more information.*

`[<none | authorized >]`

*Provides options for secondary authentication. The **none** option specifies that a backup authentication method is not used. The **authorized** option allows access without authentication. (default: **none**).*

For example, to enable the switch to perform 802.1X authentication using one or more EAP-capable RADIUS servers:

ProCurve(config)# aaa authentication port-access eap-radius
ProCurve(config)# show auth

Configuration command for EAP-RADIUS authentication.

Status and Counters - Authentication Information

Login Attempts : 3
Respect Privilege : Disabled

Access Task	Login Primary	Login Secondary	Enable Primary	Enable Secondary
Console	Local	None	Local	None
Telnet	Local	None	Local	None
Port-Access	EapRadius			
Webui	Local	None	Local	None
SSH	Local	None	Local	None
Web-Auth	ChapRadius			
MAC-Auth	ChapRadius			

802.1X (Port-Access) configured for EAP-RADIUS authentication.

Figure 10-6. Example of 802.1X (Port-Access) Authentication

4. Enter the RADIUS Host IP Address(es)

If you select either **eap-radius** or **chap-radius** for the authentication method, configure the switch to use 1, 2, or 3 RADIUS servers for authentication. The following syntax shows the basic commands. For coverage of all commands related to RADIUS server configuration, refer to chapter 5, “RADIUS Authentication, Authorization, and Accounting”.

Syntax: radius host < ip-address > [oobm]

Adds a server to the RADIUS configuration.

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the RADIUS traffic will go through the out-of-band management (OOBM) port.*

[key < server-specific key-string >]

Optional. Specifies an encryption key for use during authentication (or accounting) sessions with the specified server. This key must match the key used on the RADIUS server. Use this option only if the specified server requires a different key than configured for the global encryption key.

Syntax: radius-server key < global key-string >

Specifies the global encryption key the switch uses for sessions with servers for which the switch does not have a server-specific key. This key is optional if all RADIUS server addresses configured in the switch include a server-specific encryption key.

5. Enable 802.1X Authentication on the Switch

After configuring 802.1X authentication as described in the preceding four sections, activate it with this command:

Syntax: aaa port-access authenticator active

Activates 802.1X port-access on ports you have configured as authenticators.

6. Optional: Reset Authenticator Operation

While 802.1X authentication is operating, you can use the following **aaa port-access authenticator** commands to reset 802.1X authentication and statistics on specified ports.

Syntax: aaa port-access authenticator < *port-list* >

[initialize]

*On the specified ports, blocks inbound and outbound traffic and restarts the 802.1X authentication process. This happens only on ports configured with **control auto** and actively operating as 802.1X authenticators.*

[reauthenticate]

On the specified ports, forces reauthentication (unless the authenticator is in “HELD” state).

[clear-statistics]

On the specified ports, clears authenticator statistics counters.

7. Optional: Configure 802.1X Controlled Directions

After you enable 802.1X authentication on specified ports, you can use the **aaa port-access controlled-directions** command to configure how a port transmits traffic before it successfully authenticates a client and enters the authenticated state.

As documented in the IEEE 802.1X standard, an 802.1X-aware port that is unauthenticated can control traffic in either of the following ways:

- In both ingress and egress directions by disabling both the reception of incoming frames and transmission of outgoing frames
- Only in the ingress direction by disabling only the reception of incoming frames.

Prerequisite. As documented in the IEEE 802.1X standard, the disabling of incoming traffic and transmission of outgoing traffic on an 802.1X-aware egress port in an unauthenticated state (using the **aaa port-access controlled-directions in** command) is supported only if:

- The port is configured as an edge port in the network using the **spanning-tree edge-port** command.

- The 802.1s Multiple Spanning Tree Protocol (MSTP) or 802.1w Rapid Spanning Tree Protocol (RSTP) is enabled on the switch. MSTP and RSTP improve resource utilization while maintaining a loop-free network.

For information on how to configure the prerequisites for using the **aaa port-access controlled-directions in** command, see Chapter 4, “Multiple Instance Spanning-Tree Operation” in the *Advanced Traffic Management Guide*.

Syntax: aaa port-access <port-list> controlled-directions <both | in>

both (default): *Incoming and outgoing traffic is blocked on an 802.1X-aware port before authentication occurs.*

in: *Incoming traffic is blocked on an 802.1X-aware port before authentication occurs. Outgoing traffic with unknown destination addresses is flooded on unauthenticated 802.1X-aware ports.*

Wake-on-LAN Traffic

The Wake-on-LAN feature is used by network administrators to remotely power on a sleeping workstation (for example, during early morning hours to perform routine maintenance operations, such as patch management and software updates).

The **aaa port-access controlled-direction in** command allows Wake-on-LAN traffic to be transmitted on an 802.1X-aware egress port that has not yet transitioned to the 802.1X authenticated state; the **controlled-direction both** setting prevents Wake-on-LAN traffic to be transmitted on an 802.1X-aware egress port until authentication occurs.

Note

Although the **controlled-direction in** setting allows Wake-on-LAN traffic to traverse the switch through unauthenticated 802.1X-aware egress ports, it does not guarantee that the Wake-on-LAN packets will arrive at their destination. For example, firewall rules on other network devices and VLAN rules may prevent these packets from traversing the network.

Operating Notes

- Using the **aaa port-access controlled-directions in** command, you can enable the transmission of Wake-on-LAN traffic on unauthenticated egress ports that are configured for any of the following port-based security features:
 - 802.1X authentication
 - MAC authentication
 - Web authentication

Because a port can be configured for more than one type of authentication to protect the switch from unauthorized access, the last setting you configure with the **aaa port-access controlled-directions** command is applied to all authentication methods configured on the switch.

For information about how to configure and use MAC and Web authentication, refer to chapter 3, “Web and MAC Authentication”.

- To display the currently configured 802.1X Controlled Directions value, enter the **show port-access authenticator config** command as shown in Figure 10-11.
- When an 802.1X-authenticated port is configured with the **controlled-directions in** setting, eavesdrop prevention is not supported on the port.

Example: Configuring 802.1X Controlled Directions

The following example shows how to enable the transmission of Wake-on-LAN traffic in the egress direction on an 802.1X-aware port before it transitions to the 802.1X authenticated state and successfully authenticates a client device.

```
ProCurve(config)# aaa port-access authenticator a10
ProCurve(config)# aaa authentication port-access eap-radius
ProCurve(config)# aaa port-access authenticator active
ProCurve(config)# aaa port-access a10 controlled-directions in
```

Figure 10-7. Example of Configuring 802.1X Controlled Directions

802.1X Open VLAN Mode

802.1X Authentication Commands	page 10-18
802.1X Supplicant Commands	page 10-49
802.1X Open VLAN Mode Commands	
[no] aaa port-access authenticator < <i>port-list</i> >	page 10-43
[auth-vid < <i>vlan-id</i> >]	
[unauth-vid < <i>vlan-id</i> >]	
802.1X-Related Show Commands	page 10-51
RADIUS server configuration	pages 10-25

Introduction

This section describes how to use the 802.1X Open VLAN mode to provide a path for clients that need to acquire 802.1X supplicant software before proceeding with the authentication process. The Open VLAN mode involves options for configuring unauthorized-client and authorized-client VLANs on ports configured as 802.1X authenticators.

Configuring the 802.1X Open VLAN mode on a port changes how the port responds when it detects a new client. In earlier releases, a “friendly” client computer not running 802.1X supplicant software could not be authenticated on a port protected by 802.1X access security. As a result, the port would become blocked and the client could not access the network. This prevented the client from:

- Acquiring IP addressing from a DHCP server
- Downloading the 802.1X supplicant software necessary for an authentication session

The 802.1X Open VLAN mode solves this problem by temporarily suspending the port’s static VLAN memberships and placing the port in a designated *Unauthorized-Client VLAN* (sometimes termed a *guest VLAN*). In this state the client can proceed with initialization services, such as acquiring IP addressing and 802.1X client software, and starting the authentication process.

Note

On ports configured to allow multiple sessions using 802.1X user-based access control, all clients must use the same untagged VLAN. On a given port where there are no currently active, authenticated clients, the first *authenticated* client determines the untagged VLAN in which the port will operate for all subsequent, overlapping client sessions.

If the switch operates in an environment where some valid clients will not be running 802.1X supplicant software and need to download it from your network. Then, because such clients would need to use the Unauthorized-Client VLAN and authenticated clients would be using a different VLAN (for security reasons), allowing multiple clients on an 802.1X port can result in blocking some or all clients needing to use the Unauthorized-Client VLAN.

On ports configured for port-based 802.1X access control, if multiple clients try to authenticate on the same port, the most recently authenticated client determines the untagged VLAN membership for that port. Clients that connect without trying to authenticate will have access to the untagged VLAN membership that is currently assigned to the port.

VLAN Membership Priorities

Following client authentication, an 802.1X port resumes membership in any tagged VLANs for which it is already assigned in the switch configuration. The port also becomes an untagged member of one VLAN according to the following order of options:

- a. **1st Priority:** The port joins a VLAN to which it has been assigned by a RADIUS server during client authentication.
- b. **2nd Priority:** If RADIUS authentication does not include assigning the port to a VLAN, then the switch assigns the port to the VLAN entered in the port's 802.1X configuration as an *Authorized-Client* VLAN, if configured.
- c. **3rd Priority:** If the port does not have an Authorized-Client VLAN configured, but does have a static, untagged VLAN membership in its configuration, then the switch assigns the port to this VLAN.

A port assigned to a VLAN by an Authorized-Client VLAN configuration (or a RADIUS server) will be an untagged member of the VLAN for the duration of the authenticated session. This applies even if the port is also configured in the switch as a tagged member of the same VLAN.

Note

After client authentication, the port resumes membership in any tagged VLANs for which it is configured. If the port is a tagged member of a VLAN used for 1 or 2 listed above, then it also operates as an untagged member of that VLAN while the client is connected. When the client disconnects, the port reverts to tagged membership in the VLAN.

Use Models for 802.1X Open VLAN Modes

You can apply the 802.1X Open VLAN mode in more than one way. Depending on your use, you will need to create one or two static VLANs on the switch for *exclusive* use by per-port 802.1X Open VLAN mode authentication:

- **Unauthorized-Client VLAN:** Configure this VLAN when unauthenticated, friendly clients will need access to some services before being authenticated or instead of being authenticated.
- **Authorized-Client VLAN:** Configure this VLAN for authenticated clients when the port is not statically configured as an untagged member of a VLAN you want clients to use, or when the port is statically configured as an untagged member of a VLAN you do not want clients to use. (A port can be configured as untagged on only one port-based VLAN. When an Authorized-Client VLAN is configured, it will always be untagged and will block the port from using a statically configured, untagged membership in another VLAN.) Note that after client authentication, the port returns to membership in any tagged VLANs for which it is configured. See the "Note", above.

Table 10-2. 802.1X Open VLAN Mode Options

802.1X Per-Port Configuration	Port Response
No Open VLAN mode:	The port automatically blocks a client that cannot initiate an authentication session.
Open VLAN mode with both of the following configured:	
Unauthorized-Client VLAN	<ul style="list-style-type: none">• When the port detects a client without 802.1X supplicant capability, it automatically becomes an untagged member of this VLAN. If you previously configured the port as a static, tagged member of the VLAN, membership temporarily changes to untagged while the client remains unauthenticated.• If the port already has a statically configured, untagged membership in another VLAN, then the port temporarily closes access to this other VLAN while in the Unauthorized-Client VLAN.• To limit security risks, the network services and access available on the Unauthorized-Client VLAN should include only what a client needs to enable an authentication session. If the port is statically configured as a tagged member of any other VLANs, access to these VLANs is blocked while the port is a member of the Unauthorized-Client VLAN. <p>Note for a Port Configured To Allow Multiple Client Sessions: If any previously authenticated clients are using a port assigned to a VLAN other than the Unauthorized-Client VLAN, then a later client that is not running 802.1X supplicant software is blocked on the port until all other, authenticated clients on the port have disconnected.</p>

802.1X Per-Port Configuration	Port Response
Authorized-Client VLAN	<ul style="list-style-type: none">After client authentication, the port drops membership in the Unauthorized-Client VLAN and becomes an untagged member of this VLAN. <p>Notes: If the client is running an 802.1X supplicant application when the authentication session begins, and is able to authenticate itself before the switch assigns the port to the Unauthorized-Client VLAN, then the port does not become a member of the Unauthorized-Client VLAN. On the switches covered in this guide, you can use the unauth-period command—page 10-23—to delay moving the port into the Unauthorized-Client VLAN.</p> <p>If RADIUS authentication assigns a VLAN and there are no other authenticated clients on the port, then the port becomes a member of the RADIUS-assigned VLAN—instead of the Authorized-Client VLAN—while the client is connected.</p> <ul style="list-style-type: none">If the port is statically configured as a tagged member of a VLAN, and this VLAN is used as the Authorized-Client VLAN, then the port temporarily becomes an untagged member of this VLAN when the client becomes authenticated.If the port is statically configured as a tagged member of a VLAN, the port returns to tagged membership in this VLAN upon successful authentication. This happens even if the RADIUS server assigns the port to another, authorized VLAN. If the port is already configured as a tagged member of a VLAN that RADIUS assigns as an authorized VLAN, then the port becomes an untagged member of that VLAN for the duration of the client connection.

802.1X Per-Port Configuration	Port Response
Open VLAN Mode with Only an Unauthorized-Client VLAN Configured:	
	<ul style="list-style-type: none">When the port detects a client, it automatically becomes an untagged member of this VLAN. To limit security risks, the network services and access available on this VLAN should include only what a client needs to enable an authentication session. If the port is statically configured as an untagged member of another VLAN, the switch temporarily removes the port from membership in this other VLAN while membership in the Unauthorized-Client VLAN exists.After the client is authenticated, and if the port is statically configured as an untagged member of another VLAN, the port's access to this other VLAN is restored. <p>Note: If RADIUS authentication assigns the port to a VLAN, this assignment overrides any statically configured, untagged VLAN membership on the port (while the client is connected).</p> <ul style="list-style-type: none">If the port is statically configured as a tagged member of a VLAN, the port returns to tagged membership in this VLAN upon successful client authentication. This happens even if the RADIUS server assigns the port to another, authorized VLAN. Note that if the port is already configured as a tagged member of a VLAN that RADIUS assigns as an authorized VLAN, then the port becomes an <i>untagged</i> member of that VLAN for the duration of the client connection. <p>Note for a Port Configured To Allow Multiple Client Sessions: If any previously authenticated clients are using a port assigned to a VLAN other than the Unauthorized-Client VLAN (such as a RADIUS-assigned VLAN), then a later client that is not running 802.1X supplicant software is blocked on the port until all other, authenticated clients on the port have disconnected. Refer to figure 10-1 on page 10-11.</p>

802.1X Per-Port Configuration	Port Response
Open VLAN Mode with Only an Authorized-Client VLAN Configured:	
	<ul style="list-style-type: none">• Port automatically blocks a client that cannot initiate an authentication session.• If the client successfully completes an authentication session, the port becomes an untagged member of this VLAN.• If the port is statically configured as a tagged member of any other VLAN, the port returns to tagged membership in this VLAN upon successful client authentication. This happens even if the RADIUS server assigns the port to another, authorized VLAN. If the port is already configured as a tagged member of a VLAN that RADIUS assigns as an authorized VLAN, then the port becomes an <i>untagged</i> member of that VLAN for the duration of the client connection.
	Note: An authorized-client VLAN configuration can be overridden by a RADIUS authentication that assigns a VLAN. (Refer to figure 10-1 on page 10-11.)

Operating Rules for Authorized-Client and Unauthorized-Client VLANs

Condition	Rule
Static VLANs used as <i>Authorized-Client</i> or <i>Unauthorized-Client</i> VLANs	These must be configured on the switch before you configure an 802.1X authenticator port to use them. (Use the vlan < vlan-id > command or the VLAN Menu screen in the Menu interface.)
VLAN Assignment Received from a RADIUS Server	If the RADIUS server specifies a VLAN for an authenticated supplicant connected to an 802.1X authenticator port, this VLAN assignment overrides any Authorized-Client VLAN assignment configured on the authenticator port. This is because membership in both VLANs is untagged, and the switch allows only one untagged, port-based VLAN membership per-port. For example, suppose you configured port A4 to place authenticated supplicants in VLAN 20. If a RADIUS server authenticates supplicant "A" and assigns this supplicant to VLAN 50, then the port can access VLAN 50 as an untagged member while the client session is running. When the client disconnects from the port, then the port drops these assignments and uses the untagged VLAN memberships for which it is statically configured. (After client authentication, the port resumes any tagged VLAN memberships for which it is already configured. For details, refer to the Note on page 10-31.)
Temporary VLAN Membership During a Client Session	<ul style="list-style-type: none">• Port membership in a VLAN assigned to operate as the Unauthorized-Client VLAN is temporary, and ends when the client receives authentication or the client disconnects from the port, whichever is first. In the case of the multiple clients allowed on switches covered in this guide, the first client to authenticate determines the untagged VLAN membership for the port until all clients have disconnected. Any other clients that cannot operate in that VLAN are blocked at that point.• Port membership in a VLAN assigned to operate as the Authorized-Client VLAN ends when the client disconnects from the port. If a VLAN assignment from a RADIUS server is used instead, the same rule applies. In the case of the multiple clients allowed on switches, the port maintains the same VLAN as long as there is any authenticated client using the VLAN. When the last client disconnects, then the port reverts to only the VLAN(s) for which it is statically configured as a member.

Condition	Rule
Effect of Unauthorized-Client VLAN session on untagged port VLAN membership	<ul style="list-style-type: none"> When an unauthenticated client connects to a port that is already configured with a static, untagged VLAN, the switch temporarily moves the port to the Unauthorized-Client VLAN (also untagged). (While the Unauthorized-Client VLAN is in use, the port does not access any other VLANs.) If the client disconnects, the port leaves the Unauthorized-Client VLAN and re-acquires membership in all the statically configured VLANs to which it belongs. If the client becomes authenticated, the port leaves the Unauthorized-Client VLAN and joins the appropriate VLAN. (Refer to “VLAN Membership Priorities” on page 10-30. In the case of the multiple clients allowed on switches, if an authenticated client is already using the port for a different VLAN, then any other unauthenticated clients needing to use the Unauthorized-Client VLAN are blocked.
Effect of Authorized-Client VLAN session on untagged port VLAN membership.	<ul style="list-style-type: none"> When a client becomes authenticated on a port that is already configured with a static, untagged VLAN, the switch temporarily moves the port to the Authorized-Client VLAN (also untagged). While the Authorized-Client VLAN is in use, the port does not have access to the statically configured, untagged VLAN. When the authenticated client disconnects, the switch removes the port from the Authorized-Client VLAN and moves it back to the untagged membership in the statically configured VLAN. (After client authentication, the port resumes any tagged VLAN memberships for which it is already configured. For details, refer to the Note on page 10-31.) <p>Note: This rule assumes:</p> <ul style="list-style-type: none"> No alternate VLAN has been assigned by a RADIUS server. No other authenticated clients are already using the port.
Multiple Authenticator Ports Using the Same Unauthorized-Client and Authorized-Client VLANs	<p>You can use the same static VLAN as the Unauthorized-Client VLAN for all 802.1X authenticator ports configured on the switch. Similarly, you can use the same static VLAN as the Authorized-Client VLAN for all 802.1X authenticator ports configured on the switch.</p> <p>Caution: Do not use the same static VLAN for both the unauthorized-client VLAN and the authorized-client VLAN. Using one VLAN for both creates a security risk by defeating the isolation of unauthenticated clients.</p>
Effect of Failed Client Authentication Attempt This rule assumes no other authenticated clients are already using the port on a different VLAN.	<p>When there is an Unauthorized-Client VLAN configured on an 802.1X authenticator port, an unauthorized client connected to the port has access only to the network resources belonging to the Unauthorized-Client VLAN. This access continues until the client disconnects from the port. (If there is no Unauthorized-Client VLAN configured on the authenticator port, the port simply blocks access for any unauthorized client.)</p>

Configuring Port-Based and User-Based Access Control (802.1X)

802.1X Open VLAN Mode

Condition	Rule
Effect of RADIUS-assigned VLAN This rule assumes no other authenticated clients are already using the port on a different VLAN.	The port joins the RADIUS-assigned VLAN as an untagged member.
IP Addressing for a Client Connected to a Port Configured for 802.x Open VLAN Mode	A client can either acquire an IP address from a DHCP server or use a manually configured IP address before connecting to the switch.
802.1X Supplicant Software for a Client Connected to a Port Configured for 802.1X Open VLAN Mode	A friendly client, without 802.1X supplicant software, connecting to an authenticator port must be able to download this software from the Unauthorized-Client VLAN before authentication can begin.
Switch with a Port Configured To Allow Multiple Authorized-Client Sessions	<p>When a new client is authenticated on a given port:</p> <ul style="list-style-type: none">• If no other clients are authenticated on that port, then the port joins one VLAN in the following order of precedence:<ol style="list-style-type: none">a. A RADIUS-assigned VLAN, if configured.b. An Authenticated-Client VLAN, if configured.c. A static, port-based VLAN to which the port belongs as an untagged member.d. Any VLAN(s) to which the port is configured as a tagged member (provided that the client can operate in that VLAN).• If another client is already authenticated on the port, then the port is already assigned to a VLAN for the previously-existing client session, and the new client must operate in this same VLAN, regardless of other factors. (This means that a client without 802.1X client authentication software cannot access a configured, Unauthenticated-Client VLAN if another, authenticated client is already using the port.)

Condition	Rule
Note: Limitation on Using an Unauthorized-Client VLAN on an 802.1X Port Configured to Allow Multiple-Client Access	You can optionally enable switches to allow up to 32 clients per-port. The Unauthorized-Client VLAN feature can operate on an 802.1X-configured port regardless of how many clients the port is configured to support. However, all clients on the same port must operate through the same untagged VLAN membership. This means that any client accessing a given port must be able to authenticate and operate on the same VLAN as any other previously authenticated clients that are currently using the port. Thus, an Unauthorized-Client VLAN configured on a switch port that allows multiple 802.1X clients cannot be used if there is already an authenticated client using the port on another VLAN. Also, a client using the Unauthenticated-Client VLAN will be blocked when another client becomes authenticated on the port. For this reason, the best utilization of the Unauthorized-Client VLAN feature is in instances where only one client is allowed per-port. Otherwise, unauthenticated clients are subject to being blocked at any time by authenticated clients using a different VLAN. (Using the same VLAN for authenticated and unauthenticated clients can create a security risk and is not recommended.)

Note	If you use the same VLAN as the Unauthorized-Client VLAN for all authenticator ports, unauthenticated clients on different ports can communicate with each other.
-------------	---

Setting Up and Configuring 802.1X Open VLAN Mode

Preparation. This section assumes use of both the Unauthorized-Client and Authorized-Client VLANs. Refer to Table 10-2 on page 10-32 for other options.

Before you configure the 802.1X Open VLAN mode on a port:

- Statically configure an “Unauthorized-Client VLAN” in the switch. The only ports that should belong to this VLAN are ports offering services and access you want available to unauthenticated clients. (802.1X authenticator ports do not have to be members of this VLAN.)

Caution

Do not allow any port memberships or network services on this VLAN that would pose a security risk if exposed to an unauthorized client.

- Statically configure an Authorized-Client VLAN in the switch. The only ports that should belong to this VLAN are ports offering services and access you want available to authenticated clients. 802.1X authenticator ports do not have to be members of this VLAN.

Note that if an 802.1X authenticator port is an untagged member of another VLAN, the port’s access to that other VLAN will be temporarily removed while an authenticated client is connected to the port. For example, if:

- i. Port A5 is an untagged member of VLAN 1 (the default VLAN).
- ii. You configure port A5 as an 802.1X authenticator port.
- iii. You configure port A5 to use an Authorized-Client VLAN.

Then, if a client connects to port A5 and is authenticated, port A5 becomes an untagged member of the Authorized-Client VLAN and is temporarily suspended from membership in the default VLAN.

- If you expect friendly clients to connect without having 802.1X supplicant software running, provide a server on the Unauthorized-Client VLAN for downloading 802.1X supplicant software to the client, and a procedure by which the client initiates the download.
- A client must either have a valid IP address configured before connecting to the switch, or download one through the Unauthorized-Client VLAN from a DHCP server. In the latter case, you will need to provide DHCP services on the Unauthorized-Client VLAN.
- Ensure that the switch is connected to a RADIUS server configured to support authentication requests from clients using ports configured as 802.1X authenticators. (The RADIUS server should not be on the Unauthorized-Client VLAN.)

Note that as an alternative, you can configure the switch to use local password authentication instead of RADIUS authentication. However, this is less desirable because it means that all clients use the same passwords and have the same access privileges. Also, you must use 802.1X supplicant software that supports the use of local switch passwords.

Caution

Ensure that you do not introduce a security risk by allowing Unauthorized-Client VLAN access to network services or resources that could be compromised by an unauthorized client.

Configuring General 802.1X Operation: These steps enable 802.1X authentication, and must be done before configuring 802.1X VLAN operation.

1. Enable 802.1X authentication on the individual ports you want to serve as authenticators. (The switch automatically disables LACP on the ports on which you enable 802.1X.) On the ports you will use as authenticators with VLAN operation, ensure that the port-control parameter is set to **auto** (the default). (Refer to “1. Enable 802.1X Authentication on Selected Ports” on page 10-19.) This setting requires a client to support 802.1X authentication (with 802.1X supplicant operation) and to provide valid credentials to get network access.

Syntax: `aaa port-access authenticator < port-list > control auto`

Activates 802.1X port-access on ports you have configured as authenticators.

2. Configure the 802.1X authentication type. Options include:

Syntax: `aaa authentication port-access < local | eap-radius | chap-radius >`

Determines the type of RADIUS authentication to use.

local: *Use the switch's local username and password for supplicant authentication (the default).*

eap-radius *Use EAP-RADIUS authentication. (Refer to the documentation for your RADIUS server.*

chap-radius *Use CHAP-RADIUS (MD5) authentication. (Refer to the documentation for your RADIUS server software.)*

3. If you selected either **eap-radius** or **chap-radius** for step 2, use the **radius host** command to configure up to three RADIUS server IP address(es) on the switch.

Syntax: radius host < ip-address > [oobm]

Adds a server to the RADIUS configuration.

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the RADIUS traffic will go through the out-of-band management (OOBM) port.*

[key < server-specific key-string >]

Optional. Specifies an encryption key for use with the specified server. This key must match the key used on the RADIUS server. Use this option only if the specified server requires a different key than configured for the global encryption key.

Syntax: radius-server key < global key-string >

Specifies the global encryption key the switch uses for sessions with servers for which the switch does not have a server-specific key. This key is optional if all RADIUS server addresses configured in the switch include a server-specific encryption key.

4. Activate authentication on the switch.

Syntax: aaa port-access authenticator active

Activates 802.1X port-access on ports you have configured as authenticators.

5. Test both the authorized and unauthorized access to your system to ensure that the 802.1X authentication works properly on the ports you have configured for port-access.

Note

If you want to implement the optional port-security feature on the switch, you should first ensure that the ports you have configured as 802.1X authenticators operate as expected. Then refer to “Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1X-Authenticated Devices” on page 10-45.

After you complete steps 1 and 2, the configured ports are enabled for 802.1X authentication (without VLAN operation), and you are ready to configure VLAN Operation.

Configuring 802.1X Open VLAN Mode. Use these commands to actually configure Open VLAN mode. For a listing of the steps needed to prepare the switch for using Open VLAN mode, refer to “Preparation” on page 10-40.

Syntax: aaa port-access authenticator < port-list >

 [auth-vid < vlan-id >]

Configures an existing, static VLAN to be the Authorized-Client VLAN.

 [< unauth-vid < vlan-id >]

Configures an existing, static VLAN to be the Unauthorized-Client VLAN.

For example, suppose you want to configure 802.1X port-access with Open VLAN mode on ports A10 - A20 and:

- These two static VLANs already exist on the switch:
 - Unauthorized, VID = 80
 - Authorized, VID = 81
- Your RADIUS server has an IP address of 10.28.127.101. The server uses **rad4all** as a server-specific key string. The server is connected to a port on the Default VLAN.
- The switch's default VLAN is already configured with an IP address of 10.28.127.100 and a network mask of 255.255.255.0

```
ProCurve(config)# aaa authentication port-access eap-radius
```

Configures the switch for 802.1X authentication using an EAP-RADIUS server.

```
ProCurve(config)# aaa port-access authenticator a10-a20
```

Configures ports A10 - A20 as 802.1 authenticator ports.

```
ProCurve(config)# radius host 10.28.127.101 key rad4all
```

Configures the switch to look for a RADIUS server with an IP address of 10.28.127.101 and an encryption key of rad4all.

```
ProCurve(config)# aaa port-access authenticator e a10-a20 unauth-vid 80
```

Configures ports A10 - A20 to use VLAN 80 as the Unauthorized-Client VLAN.

```
ProCurve(config)# aaa port-access authenticator e a10-a20 auth-vid 81
```

Configures ports A10 - A20 to use VLAN 81 as the Authorized-Client VLAN.

```
ProCurve(config)# aaa port-access authenticator active
```

Activates 802.1X port-access on ports you have configured as authenticators.

Inspecting 802.1X Open VLAN Mode Operation. For information and an example on viewing current Open VLAN mode operation, refer to “Viewing 802.1X Open VLAN Mode Status” on page 10-60.

802.1X Open VLAN Operating Notes

- Although you can configure Open VLAN mode to use the same VLAN for both the Unauthorized-Client VLAN and the Authorized-Client VLAN, this is *not* recommended. Using the same VLAN for both purposes allows unauthenticated clients access to a VLAN intended only for authenticated clients, which poses a security breach.
- While an Unauthorized-Client VLAN is in use on a port, the switch temporarily removes the port from any other statically configured VLAN for which that port is configured as a member. Note that the Menu interface will still display the port’s statically configured VLAN(s).
- A VLAN used as the Unauthorized-Client VLAN should not allow access to resources that must be protected from unauthenticated clients.
- If a port is configured as a tagged member of VLAN “X”, then the port returns to tagged membership in VLAN “X” upon successful client authentication. This happens even if the RADIUS server assigns the port to another, authorized VLAN “Y”. Note that if RADIUS assigns VLAN “X” as an authorized VLAN, then the port becomes an *untagged* member of VLAN “X” for the duration of the client connection. (If there is no Authorized-Client or RADIUS-assigned VLAN, then an authenticated client without tagged VLAN capability can access only a statically configured, untagged VLAN on that port.)
- When a client’s authentication attempt on an Unauthorized-Client VLAN fails, the port remains a member of the Unauthorized-Client VLAN until the client disconnects from the port.
- During an authentication session on a port in 802.1X Open VLAN mode, if RADIUS specifies membership in an untagged VLAN, this assignment overrides port membership in the Authorized-Client VLAN. If there is no Authorized-Client VLAN configured, then the RADIUS assignment overrides any untagged VLAN for which the port is statically configured.
- If the only authenticated client on a port loses authentication during a session in 802.1X Open VLAN mode, the port VLAN membership reverts back to the Unauthorized-Client VLAN. If there is no Unauthorized-Client VLAN configured, then the client loses access to the port until it can reauthenticate itself. If there are multiple clients authenticated on the port, if one client loses access and attempts to re-authenticate, that client will be handled as a new client on the port.

- The first client to authenticate on a port configured to support multiple clients will determine the port’s VLAN membership for any subsequent clients that authenticate while an active session is already in effect.

Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1X-Authenticated Devices

If 802.1X authentication is disabled on a port or set to **authorized** (Force Authorize), the port can allow access to a non-authenticated client. Port-Security operates with 802.1X authentication only if the selected ports are configured as 802.1X *with* the **control** mode in the port-access authenticator command set to **auto** (the default setting). For example, if port A10 was at a non-default 802.1X setting and you wanted to configure it to support the port-security option, you would use the following **aaa port-access** command:

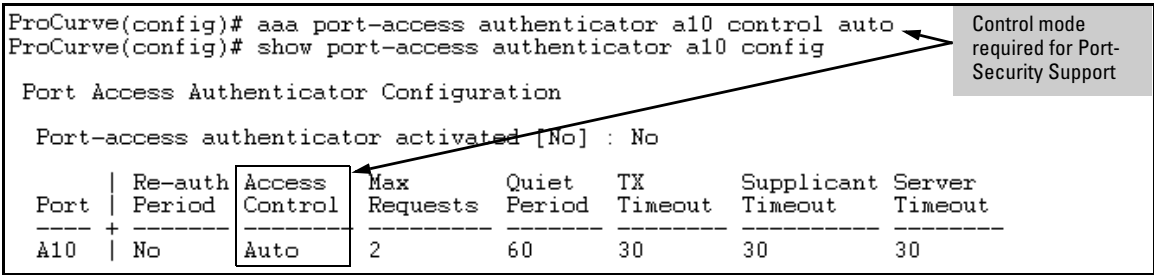


Figure 10-8. Port-Access Support for Port-Security Operation

Port-Security

Note

If 802.1X port-access is configured on a given port, then port-security **learn-mode** for that port must be set to either **continuous** (the default) or **port-access**.

In addition to the above, to use port-security on an authenticator port (chapter 11), use the per-port **client-limit** option to control how many MAC addresses of 802.1X-authenticated devices the port is allowed to learn. (Using **client-limit** sets 802.1X to user-based operation on the specified ports.) When this limit is reached, no further devices can be authenticated until a currently authenticated device disconnects and the current delay period or logoff period has expired.

Configure the port access type.

Syntax: aaa port-access auth < port-list > client-limit < 1 - 32 >

Configures user-based 802.1X authentication on the specified ports and sets the number of authenticated devices the port is allowed to learn. For more on this command, refer to “Configuring Switch Ports as 802.1X Authenticators” on page 10-18.)

— Or —

no aaa port-access auth < port-list > client-limit

Configures port-based 802.1X authentication on the specified ports, which opens the port. (Refer to “User Authentication Methods” on page 10-4.)

Configuring Switch Ports To Operate As Supplicants for 802.1X Connections to Other Switches

802.1X Authentication Commands	page 10-18
802.1X Supplicant Commands	
[no] aaa port-access < supplicant < [ethernet] < <i>port-list</i> >	page 10-49
[auth-timeout held-period start-period max-start initialize identity secret clear-statistics]	page 10-49
802.1X-Related Show Commands	page 10-51
RADIUS server configuration	pages 10-25

A switch port can operate as a supplicant in a connection to a port on another 802.1X-aware switch to provide security on links between 802.1X-aware switches. (A port can operate as both an authenticator and a supplicant.)

Example

Suppose that you want to connect two switches, where:

- Switch “A” has port A1 configured for 802.1X supplicant operation.
- You want to connect port A1 on switch “A” to port B5 on switch “B”.




Figure 10-9. Example of Supplicant Operation

1. When port A1 on switch “A” is first connected to a port on switch “B”, or if the ports are already connected and either switch reboots, port A1 begins sending start packets to port B5 on switch “B”.

Configuring Port-Based and User-Based Access Control (802.1X)

Configuring Switch Ports To Operate As Supplicants for 802.1X Connections to Other Switches

- If, after the supplicant port sends the configured number of start packets, it does not receive a response, it assumes that switch “B” is not 802.1X-aware, and transitions to the authenticated state. If switch “B” is operating properly and is not 802.1X-aware, then the link should begin functioning normally, but without 802.1X security.
 - If, after sending one or more start request packets, port A1 receives a request packet from port B5, then switch “B” is operating as an 802.1X authenticator. The supplicant port then sends a response/ID packet. If switch “B” is configured for RADIUS authentication, it forwards this request to a RADIUS server. If switch “B” is configured for Local 802.1X authentication, the authenticator compares the switch “A” response to its local username and password.
2. The RADIUS server then responds with an MD5 access challenge that switch “B” forwards to port A1 on switch “A”.
 3. Port A1 replies with an MD5 hash response based on its username and password or other unique credentials. Switch “B” forwards this response to the RADIUS server.
 4. The RADIUS server then analyzes the response and sends either a “success” or “failure” packet back through switch “B” to port A1.
 - A “success” response unblocks port B5 to normal traffic from port A1.
 - A “failure” response continues the block on port B5 and causes port A1 to wait for the “held-time” period before trying again to achieve authentication through port B5.

Supplicant Port Configuration

Enabling a Switch Port as a Supplicant. You can configure a switch port as a supplicant for a point-to-point link to an 802.1X-aware port on another switch. *Configure the port as a supplicant before configuring any supplicant-related parameters.*

Syntax: [no] aaa port-access supplicant [ethernet] < port-list >

Configures a port as a supplicant with either the default supplicant settings or any previously configured supplicant settings, whichever is most recent. The “no” form of the command disables supplicant operation on the specified ports.

Configuring a Supplicant Switch Port. You must enable supplicant operation on a port before changing the supplicant configuration. This means you must execute the supplicant command once without any other parameters, then execute it again with a supplicant parameter you want to configure. If the intended authenticator port uses RADIUS authentication, then use the **identity** and **secret** options to configure the RADIUS-expected credentials on the supplicant port. If the intended authenticator port uses Local 802.1X authentication, then use the **identity** and **secret** options to configure the authenticator switch’s local username and password on the supplicant port.

Syntax: aaa port-access supplicant [ethernet] < port-list >

*To enable supplicant operation on the designated ports, execute this command without any other parameters. After doing this, you can use the command again with the following parameters to configure supplicant operation. (Use one instance of the command for each parameter you want to configure. The **no** form disables supplicant operation on the designated port(s).*

[identity < username >]

Sets the username and password to pass to the authenticator port when a challenge-request packet is received from the authenticator port due to an authentication request. If the intended authenticator port is configured for RADIUS authentication, then < username > and < password > must be the username and password expected by the RADIUS server. If the intended authenticator port is configured for Local authentication, then < username > and < password > must be the username and password configured on the Authenticator switch. (Default: Null.)

aaa port-access supplicant [ethernet] < port-list > (**Syntax Continued**)

[secret]

Enter secret: < password >

Repeat secret: < password >

Sets the secret password to be used by the port supplicant when an MD5 authentication request is received from an authenticator. The switch prompts you to enter the secret password after the command is invoked.

[auth-timeout < 1 - 300 >]

*Sets the delay period the port waits to receive a challenge from the authenticator. If the request times out, the port sends another request, up to the number of attempts specified by the **max-start** parameter. (Default: 30 seconds).*

[max-start < 1 - 10 >]

Defines the maximum number of times the supplicant port requests authentication. See step 1 on page 10-47 for a description of how the port reacts to the authenticator response. (Default: 3).

[held-period < 0 - 65535 >]

Sets the time period the supplicant port waits after an active 802.1X session fails before trying to re-acquire the authenticator port. (Default: 60 seconds)

[start-period < 1 - 300 >]

*Sets the delay between Start packet retransmissions. That is, after a supplicant sends a start packet, it waits during the **start-period** for a response. If no response comes during the **start-period**, the supplicant sends a new start packet. The **max-start** setting (above) specifies how many start attempts are allowed in the session. (Default: 30 seconds)*

aaa port-access supplicant [ethernet] < port-list >

[initialize]

On the specified ports, blocks inbound and outbound traffic and restarts the 802.1X authentication process. Affects only ports configured as 802.1X supplicants.

[clear-statistics]

Clears and restarts the 802.1X supplicant statistics counters.

Displaying 802.1X Configuration, Statistics, and Counters

802.1X Authentication Commands	page 10-18
802.1X Supplicant Commands	page 10-47
802.1X Open VLAN Mode Commands	page 10-29
802.1X-Related Show Commands	
show port-access authenticator	page 10-53
show port-access authenticator config	page 10-53
show port-access authenticator statistics	page 10-55
show port-access authenticator session-counters	page 10-56
show port-access authenticator vlan	page 10-57
show port-access authenticator clients	page 10-58
show port-access authenticator clients <port-list> detailed	page 10-59
show port-access supplicant	page 10-64
Details of 802.1X Mode Status Listings	page 10-60
RADIUS server configuration	pages 10-25

Show Commands for Port-Access Authenticator

Syntax: show port-access authenticator [*port-list*]
 [config | statistics | session-counters | vlan | clients [detailed]]

*If you enter the **show port-access authenticator** command without an optional value, the following configuration information is displayed for all switch ports, or specified ports, that are enabled for 802.1X port-access authentication:*

- **Port -access authenticator activated:** *Are any switch ports configured to operate as 802.1X authenticators using the **aaa port-access authenticator** command? **Yes** or **No***
- **Allow RADIUS-assigned dynamic (GVRP) VLANs:** *Are RADIUS-assigned dynamic (GVRP-learned) VLANs supported for authenticated and unauthenticated client sessions on the switch? **Yes** or **No***
- **Auth Clients:** *Number of authorized clients*
- **Unauth Clients:** *Number of unauthorized clients*

Syntax: show port-access authenticator [*port-list*]
[config | statistics | session-counters | vlan | clients | clients detailed]

- **Untagged VLAN:** *VLAN ID number of the untagged VLAN used in client sessions. If the switch supports MAC-based (untagged) VLANs, **MACbased** is displayed to show that multiple untagged VLANs are configured for authentication sessions.*
- **Tagged VLANs:** *Are tagged VLANs (statically configured or RADIUS-assigned) used for authenticated clients?*
Yes or No
- **Port CoS:**
Yes - *Client-specific CoS (Class of Service) values are applied to more than one authenticated client on the port.*
No - *No client-specific CoS values are applied to any authenticated client on the port.*
*<cos-value>- Numerical value of the CoS (802.1p priority) applied to inbound traffic from one authenticated client. For client-specific per-port CoS values, enter the **show port-access web-based clients detailed** command.*
- **% In Limit:**
Yes - *Per-port rate-limiting is applied to inbound traffic from more than one authenticated client.*
No - *Per-port rate-limiting is not applied to any authenticated client session.*
<rate-limit-value>- Percentage value of a port's bandwidth applied as a rate limit to one authenticated client.

Information on ports not enabled for 802.1X port-access authentication is not displayed.

```
ProCurve(config)# show port-access authenticator
```

```
Port Access Authenticator Status
```

```
Port-access authenticator activated [No] : Yes
```

```
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : Yes
```

	Auth	Unauth		Untagged	Tagged	Port	Cntrl
Port	Clients	Clients		VLAN	VLANs	COS	Dir
----	-----	-----	+	-----	-----	-----	-----
1	1	1		4006	Yes	70000000	
2	2	0		MACbased	No	Yes	
3	4	0		1	Yes	No	
...							

Figure 10-10.Example of show port-access authenticator Command

The information displayed with the **show port-access authenticator** command for individual (config | statistics | session-counters | vlan | clients) options is described below.

Syntax: show port-access authenticator config [port-list]

Displays 802.1X port-access authenticator configuration settings, including:

- *Whether port-access authentication is enabled*
- *Whether RADIUS-assigned dynamic VLANs are supported*
- *802.1X configuration of ports that are enabled as 802.1X authenticators (For a description of each setting, refer to the syntax descriptions in “2. Reconfigure Settings for Port-Access” on page 10-21. Use the **show running** command to view the current **client-limit** configuration available for switches.)*

You can display 802.1X port-access authenticator configuration for all switch ports or specified ports.

802.1X configuration information for ports that are not enabled as 802.1X authenticators is not displayed.

```
ProCurve(config)# show port-access authenticator config

Port Access Authenticator Configuration

Port-access authenticator activated [No] : Yes
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

  Port | Re-auth Access  Max  Quiet  TX      Supplicant  Server  Cntrl
        | Period Control Reqs  Period Timeout Timeout Timeout Dir
  ----+-----
  1    | No      Auto    2    60    30      30      30    both
  2    | No      Auto    2    60    30      30      30    in
  ...
```

Figure 10-11.Example of show port-access authenticator config Command

Table 10-3. Field Descriptions of show port-access authenticator config Command Output (Figure 10-11)

Field	Description
Port-access authenticator activated	Whether 802.1X authentication is enabled or disabled on specified port(s).
Port	Port number on switch.
Re-auth Period	Period of time (in seconds) after which clients connected to the port need to be re-authenticated.
Access Control	Port's authentication mode: Auto: Network access is allowed to any connected device that supports 802.1X authentication and provides valid 802.1X credentials. Authorized: Network access is allowed to any device connected to the port, regardless of whether it meets 802.1X criteria. Unauthorized: Network access is blocked to any device connected to the port, regardless of whether the device meets 802.1X criteria.
Max reqs	Number of authentication attempts that must time-out before authentication fails and the authentication session ends.
Quiet Period	Period of time (in seconds) during which the port does not try to acquire a supplicant.
TX Timeout	Period of time (in seconds) that the port waits to retransmit the next EAPOL PDU during an authentication session.
Supplicant Timeout	Period of time (in seconds) that the switch waits for a supplicant response to an EAP request.
Server Timeout	Period of time (in seconds) that the switch waits for a server response to an authentication request.
Cntrl Dir	Directions in which flow of incoming and outgoing traffic is blocked on 802.1X-aware port that has not yet entered the authenticated state: Both: Incoming and outgoing traffic is blocked on port until authentication occurs. In: Only incoming traffic is blocked on port before authentication occurs. Outgoing traffic with unknown destination addresses is flooded on the unauthenticated 802.1X-aware port.

Syntax: show port-access authenticator statistics [*port-list*]

Displays statistical information for all switch ports or specified ports that are enabled as 802.1X authenticators, including:

- *Whether port-access authentication is enabled*
- *Whether RADIUS-assigned dynamic VLANs are supported*
- *802.1X supplicant's MAC address as determined by the content of the last EAPOL frame received on the port*
- *802.1X traffic statistics from received and transmitted packets*

802.1X configuration information for ports that are not enabled as an 802.1X authenticators is not displayed.

```
ProCurve(config)# show port-access authenticator statistics
```

Port Access Authenticator Statistics

Port-access authenticator activated [No] : Yes

Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

Port	Source MAC address	TX ReqId	TX Req	RX Start	RX Logoff	RX RespId	RX Resp	RX Errors
2	001560-b3ea48	1	0	0	0	0	0	0

Figure 10-12.Example of show port-access authenticator statistics Command

Syntax: show port-access authenticator session-counters [*port-list*]

Displays information for active 802.1X authentication sessions on all switch ports or specified ports that are enabled as 802.1X authenticators, including:

- *802.1X frames received and transmitted on each port*
 - *Duration and status of active 802.1X authentication sessions (**in-progress** or **terminated**)*
 - *User name of 802.1X supplicant included in 802.1X response packets, configured with the **aaa port-access supplicant-identity** <username> command (see page 13-49).*
- 802.1X configuration information for ports that are not enabled as an 802.1X authenticators is not displayed.*

```
ProCurve(config)# show port-access authenticator session-counters

Port Access Authenticator Session Counters

Port-access authenticator activated [No] : Yes
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

Port  Frames In  Frames Out  Session  Session
-----  -
2      45623      45623      Time(sec.)  Status      User
-----  -
                20          in-progress bert
```

Figure 10-13.Example of show port-access authenticator session-counters Command

Syntax: show port-access authenticator vlan [*port-list*]

Displays the following information on the VLANs configured for use in 802.1X port-access authentication on all switch ports, or specified ports, that are enabled as 802.1X authenticator:

- *Authentication mode used on each port, configured with the **aaa port-access authenticator control** command (see page 10-21)*
- *VLAN ID (if any) to be used for traffic from 802.1X-authenticated clients*
- *VLAN ID (if any) to be used for traffic from unauthenticated clients*

802.1X configuration information for ports that are not enabled as an 802.1X authenticators is not displayed.

```
ProCurve(config)# show port-access authenticator vlan

Port Access Authenticator VLAN Configuration

Port-access authenticator activated [No] : Yes
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

      Access   Unauth   Auth
Port Control VLAN ID  VLAN ID
----
2      Auto    0          0
```

Figure 10-14.Example of show port-access authenticator vlan Command

Syntax: show port-access authenticator clients [*port-list*]

Displays the session status, name, and address for each 802.1X port-access-authenticated client on the switch. Multiple authenticated clients may be displayed for the same port. The IP address displayed is taken from the DHCP binding table (learned through the DHCP Snooping feature).

- *If DHCP snooping is not enabled on the switch, **n/a** (not available) is displayed for a client's IP address.*
- *If an 802.1X-authenticated client uses an IPv6 address, **n/a - IPv6** is displayed.*
- *If DHCP snooping is enabled but no MAC-to-IP address binding for a client is found in the DHCP binding table, **n/a - no info** is displayed.*

ProCurve (config)# show port-access authenticator clients				
Port Access Authenticator Client Status				
Port	Client Name	MAC Address	IP Address	Session Status
1	webuser1	001321-eb8063	192.192.192.192	Authenticated
1	webuser2	001560-b3ea48	n/a - no info	Authenticating
1	webuser3	000000-111111	n/a - IPv6	Authenticating
3	webuser4	000000-111112	n/a	Authenticating

Figure 10-15. Example of show port-access authenticator clients Command Output

Syntax: show port-access authenticator clients <port-list> detailed

Displays detailed information on the status of 802.1X-authenticated client sessions on specified ports.

```
ProCurve (config)# show port-access authenticator clients 5 detailed

Port Access Authenticator Client Status Detailed

Client Base Details :
  Port           : 5
  Session Status : Open           Session Time(sec) : 999999999
  Frames In      : 999999999      Frames Out       : 999999999
  Username       : webuser1       MAC Address      : 001321-eb8063
  IP             : 2001:fe3d:ba23:cd1f:dcb1:1010:9234:4088

Access Policy Details :
  COS Map        : 70000000      In Limit %       : 87
  Untagged VLAN  : 3096          Out Limit %      : 100
  Tagged VLANs   : 1, 3, 5, 6, 334, 2066

RADIUS-ACL List :
  deny in udp from any to 10.2.8.233 CNT
    Hit Count: 0
  permit in udp from any to 10.2.8.233 CNT
    Hit Count: 0
  deny in tcp from any to 10.2.8.233 CNT
    Hit Count: 0
  permit in tcp from any to 10.2.8.233 CNT
    Hit Count: 0
  permit in tcp from any to 0.0.0.0/0 CNT
    Hit Count: 0
```

Figure 10-16. Example of show port-access authenticator clients detailed Command Output

Viewing 802.1X Open VLAN Mode Status

You can examine the switch's current VLAN status by using the **show port-access authenticator vlan** and **show port-access authenticator < port-list >** commands as illustrated in figure 10-17. Table 10-2 describes the data that these two commands display. Figure 10-18 shows related VLAN data that can help you to see how the switch is using statically configured VLANs to support 802.1X operation.

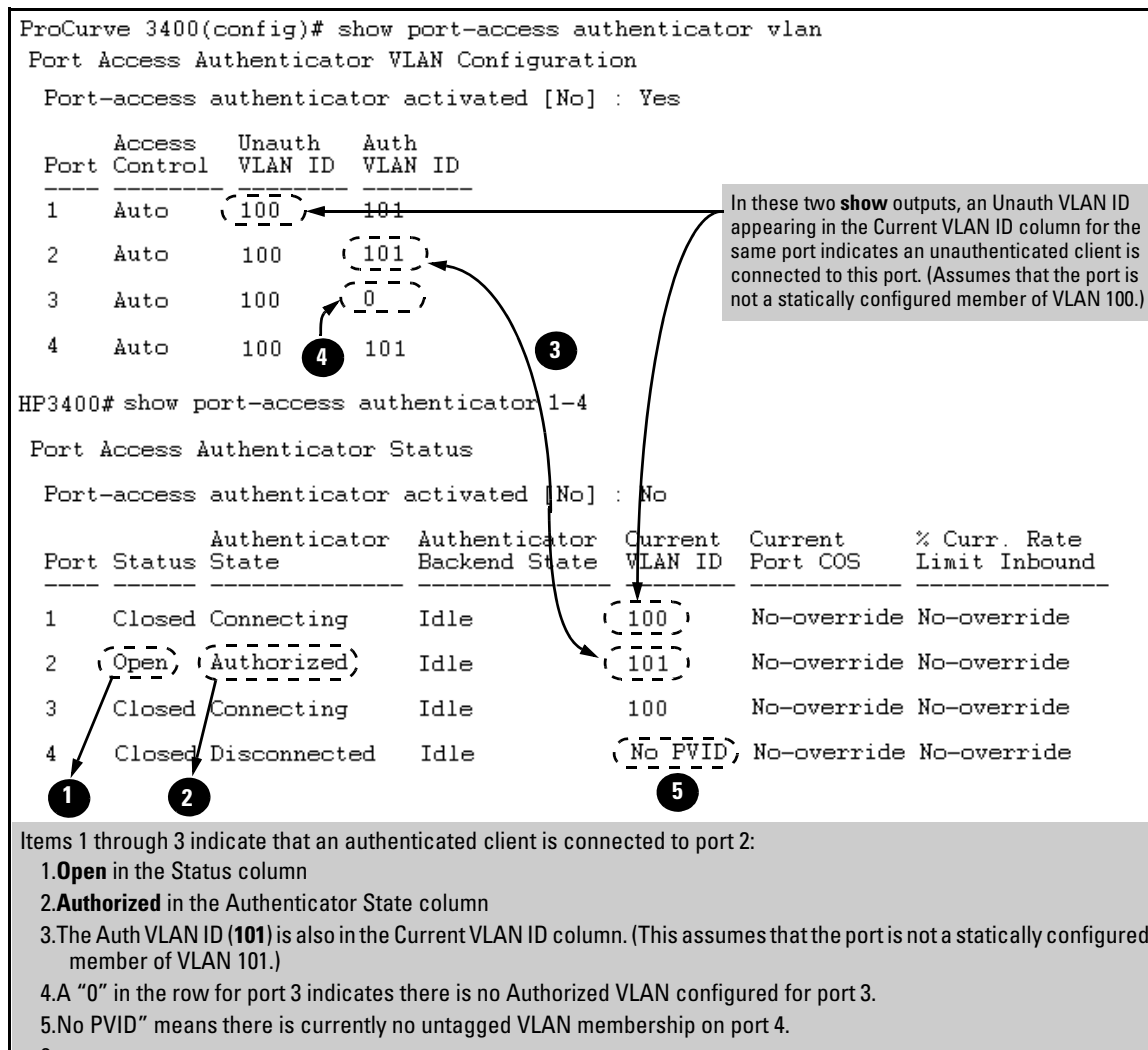


Figure 10-17. Example Showing Ports Configured for Open VLAN Mode

Thus, in the output shown in figure 10-17:

- When the **Auth VLAN ID** is configured and matches the **Current VLAN ID**, an authenticated client is connected to the port. (This assumes the port is not a statically configured member of the VLAN you are using for Auth VLAN.)
- When the **Unauth VLAN ID** is configured and matches the **Current VLAN ID**, an unauthenticated client is connected to the port. (This assumes the port is not a statically configured member of the VLAN you are using for Unauth VLAN.)

Note that because a temporary Open VLAN port assignment to either an authorized or unauthorized VLAN is an untagged VLAN membership, these assignments temporarily replace any other untagged VLAN membership that is statically configured on the port. For example, if port 12 is statically configured as an untagged member of VLAN 1, but is configured to use VLAN 25 as an authorized VLAN, then the port's membership in VLAN 1 will be temporarily suspended whenever an authenticated 802.1X client is attached to the port.

Table 10-1. Output for Determining Open VLAN Mode Status (Figure 10-17, Upper)

Status Indicator	Meaning
Access Control	
This state is controlled by the following port-access command syntax:	
ProCurve(config)# aaa port-access authenticator < port-list > control < authorized auto unauthorized >	
Auto: Configures the port to allow network access to any connected device that supports 802.1X authentication and provides valid 802.1X credentials. (This is the default authenticator setting.)	
Authorized: Configures the port for "Force Authorized", which allows access to any device connected to the port, regardless of whether it meets 802.1X criteria. (You can still configure console, Telnet, or SSH security on the port.)	
Unauthorized: Configures the port for "Force Unauthorized", which blocks access to any device connected to the port, regardless of whether the device meets 802.1X criteria.	
Unauthorized VLAN ID	<p>< <i>vlan-id</i> >: Lists the VID of the static VLAN configured as the unauthorized VLAN for the indicated port.</p> <p>0: No unauthorized VLAN has been configured for the indicated port.</p>
Authorized VLAN ID	<p>< <i>vlan-id</i> >: Lists the VID of the static VLAN configured as the authorized VLAN for the indicated port.</p> <p>0: No authorized VLAN has been configured for the indicated port.</p>

Table 10-3. Output for Determining Open VLAN Mode Status (Figure 10-17, Lower)

Status Indicator	Meaning
Status	Closed: Either no client is connected or the connected client has not received authorization through 802.1X authentication. Open: An authorized 802.1X supplicant is connected to the port.
Current VLAN ID	< vlan-id >: Lists the VID of the static, untagged VLAN to which the port currently belongs. No PVID: The port is not an untagged member of any VLAN.
Current Port CoS	<i>Refer to the section describing RADIUS support for Identity-Driven Management—IDM—in chapter 5, “RADIUS Authentication, Authorization, and Accounting” in this guide.</i>
% Curr. Rate Limit Inbound	

Syntax: show vlan < vlan-id >

Displays the port status for the selected VLAN, including an indication of which port memberships have been temporarily overridden by Open VLAN mode.

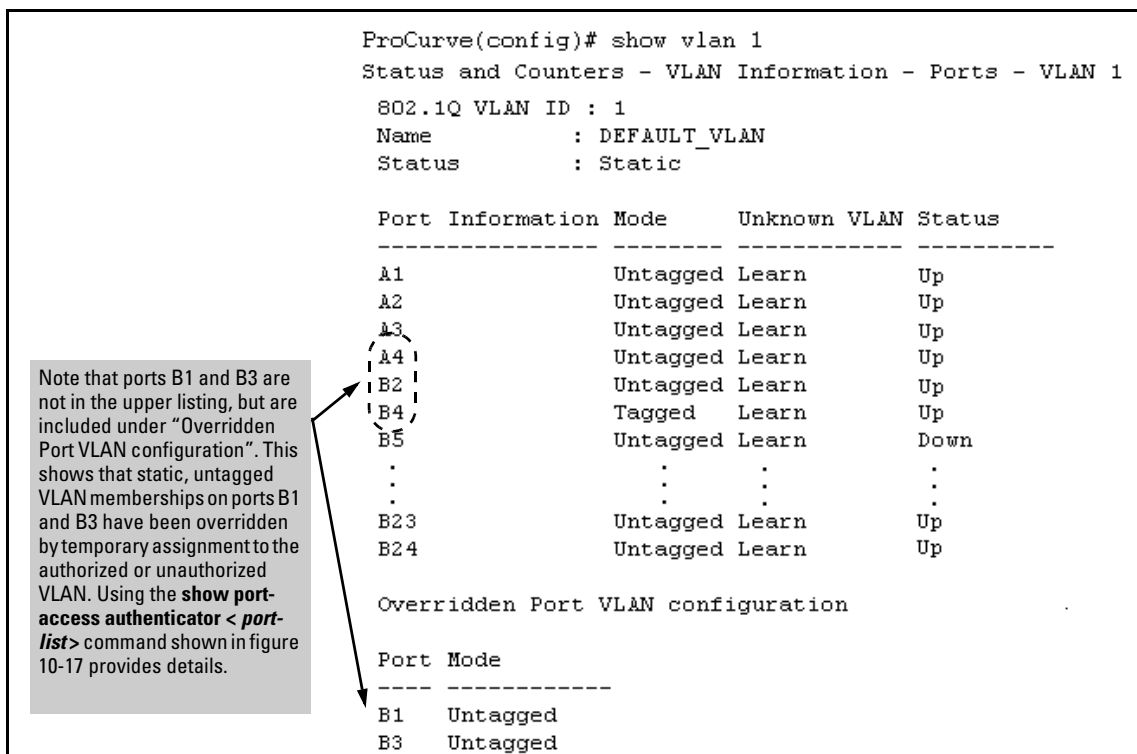


Figure 10-18. Example of Showing a VLAN with Ports Configured for Open VLAN Mode

Show Commands for Port-Access Supplicant

Syntax: show port-access supplicant [*< port-list >*] [statistics]

show port-access supplicant [*< port-list >*]

*Shows the port-access supplicant configuration (excluding the **secret** parameter) for all ports or < port-list > ports configured on the switch as supplicants. The Supplicant State can include the following:*

Connecting - Starting authentication.

Authenticated - Authentication completed (regardless of whether the attempt was successful).

Acquired - The port received a request for identification from an authenticator.

Authenticating - Authentication is in progress.

Held - Authenticator sent notice of failure. The supplicant port is waiting for the authenticator's held-period (page 10-49).

For descriptions of the supplicant parameters, refer to "Configuring a Supplicant Switch Port" on page 10-49.

show port-access supplicant [*< port-list >*] statistics

Shows the port-access statistics and source MAC address(es) for all ports or < port-list > ports configured on the switch as supplicants. See the "Note on Supplicant Statistics", below.

Note on Supplicant Statistics. For each port configured as a supplicant, **show port-access supplicant statistics < port-list >** displays the source MAC address and statistics for transactions with the authenticator device most recently detected on the port. If the link between the supplicant port and the authenticator device fails, the supplicant port continues to show data received from the connection to the most recent authenticator device until one of the following occurs:

- The supplicant port detects a different authenticator device.
- You use the **aaa port-access supplicant < port-list > clear-statistics** command to clear the statistics for the supplicant port.
- The switch reboots.

Thus, if the supplicant's link to the authenticator fails, the supplicant retains the transaction statistics it most recently received until one of the above events occurs. Also, if you move a link with an authenticator from one

supplicant port to another without clearing the statistics data from the first port, the authenticator's MAC address will appear in the supplicant statistics for both ports.

How RADIUS/802.1X Authentication Affects VLAN Operation

Static VLAN Requirement. RADIUS authentication for an 802.1X client on a given port can include a (static) VLAN requirement. (Refer to the documentation provided with your RADIUS application.) The static VLAN to which a RADIUS server assigns a client must already exist on the switch. If it does not exist or is a dynamic VLAN (created by GVRP), authentication fails. Also, for the session to proceed, the port must be an untagged member of the required VLAN. If it is not, the switch temporarily reassigns the port as described below.

If the Port Used by the Client Is Not Configured as an Untagged Member of the Required Static VLAN: When a client is authenticated on port “N”, if port “N” is not already configured as an untagged member of the static VLAN specified by the RADIUS server, then the switch temporarily assigns port “N” as an untagged member of the required VLAN (for the duration of the 802.1X session). *At the same time, if port “N” is already configured as an untagged member of another VLAN, port “N” loses access to that other VLAN for the duration of the session.* (This is because a port can be an untagged member of only one VLAN at a time.)

Using a RADIUS server to authenticate clients, you can provide port-level security protection from unauthorized network access for the following authentication methods:

- 802.1X: Port-based or client-based access control to open a port for client access after authenticating valid user credentials.
- MAC address: Authenticates a device's MAC address to grant access to the network.
- Web-browser interface: Authenticates clients for network access using a web page for user login.

Note

You can use 802.1X (port-based or client-based) authentication and either Web or MAC authentication at the same time on a port, with a maximum of 32 clients allowed on the port. (The default is one client.) Web authentication and MAC authentication are mutually exclusive on the same port. Also, you must disable LACP on ports configured for any of these authentication methods. For more information, see “Web and MAC Authentication” on page 3-1 in this guide.

VLAN Assignment on a Port

Following client authentication, VLAN configurations on a port are managed as follows when you use 802.1X, MAC, or Web authentication:

- The port resumes membership in any tagged VLANs for which it is already assigned in the switch configuration. Tagged VLAN membership allows a port to be a member of multiple VLANs simultaneously.
- The port is temporarily assigned as a member of an untagged (static or dynamic) VLAN for use during the client session according to the following order of options.
 - a. The port joins the VLAN to which it has been assigned by a RADIUS server during client authentication.
 - b. If RADIUS authentication does not include assigning the port to a VLAN, then the switch assigns the port to the authorized-client VLAN configured for the authentication method.
 - c. If the port does not have an authorized-client VLAN configured, but is configured for membership in an untagged VLAN, the switch assigns the port to this untagged VLAN.

Operating Notes

- During client authentication, a port assigned to a VLAN by a RADIUS server or an authorized-client VLAN configuration is an untagged member of the VLAN for the duration of the authenticated session. This applies even if the port is also configured in the switch as a tagged member of the same VLAN. The following restrictions apply:
 - If the port is assigned as a member of an untagged *static* VLAN, the VLAN must already be configured on the switch. If the static VLAN configuration does not exist, the authentication fails.

- If the port is assigned as a member of an untagged *dynamic* VLAN that was learned through GVRP, the dynamic VLAN configuration must exist on the switch at the time of authentication and GVRP-learned dynamic VLANs for port-access authentication must be enabled.

If the dynamic VLAN does not exist or if you have not enabled the use of a dynamic VLAN for authentication sessions on the switch, the authentication fails.

- To enable the use of a GVRP-learned (dynamic) VLAN as the untagged VLAN used in an authentication session, enter the **aaa port-access gvrp-vlans** command, as described in “Enabling the Use of GVRP-Learned Dynamic VLANs in Authentication Sessions” on page 10-71.
- Enabling the use of dynamic VLANs in an authentication session offers the following benefits:
 - You avoid the need of having static VLANs pre-configured on the switch.
 - You can centralize the administration of user accounts (including user VLAN IDs) on a RADIUS server.

For information on how to enable the switch to dynamically create 802.1Q-compliant VLANs on links to other devices using the GARP VLAN Registration Protocol (GVRP), see the chapter on “GVRP” in the *Advanced Traffic Management Guide*.

- For an authentication session to proceed, a port must be an untagged member of the (static or dynamic) VLAN assigned by the RADIUS server (or an authorized-client VLAN configuration). The port temporarily drops any current untagged VLAN membership.

If the port is not already a member of the RADIUS-assigned (static or dynamic) untagged VLAN, the switch temporarily reassigns the port as an untagged member of the required VLAN (for the duration of the session). *At the same time, if the port is already configured as an untagged member of a different VLAN, the port loses access to the other VLAN for the duration of the session.* (A port can be an untagged member of only one VLAN at a time.)

When the authentication session ends, the switch removes the temporary untagged VLAN assignment and re-activates the temporarily disabled, untagged VLAN assignment.

- If GVRP is already enabled on the switch, the temporary untagged (static or dynamic) VLAN created on the port for the authentication session is advertised as an existing VLAN.

If this temporary VLAN assignment causes the switch to disable a different untagged static or dynamic VLAN configured on the port (as described in the preceding bullet and in “Example of Untagged VLAN Assignment in a RADIUS-Based Authentication Session” on page 10-68), the disabled VLAN assignment is not advertised. When the authentication session ends, the switch:

- Removes the temporary untagged VLAN assignment and stops advertising it.
 - Re-activates and resumes advertising the temporarily disabled, untagged VLAN assignment.
- If you modify a VLAN ID configuration on a port during an 802.1X, MAC, or Web authentication session, the changes do not take effect until the session ends.
 - When a switch port is configured with RADIUS-based authentication to accept multiple 802.1X and/or MAC or Web authentication client sessions, all authenticated clients must use the same port-based, untagged VLAN membership assigned for the earliest, currently active client session.

Therefore, on a port where one or more authenticated client sessions are already running, all such clients are on the same untagged VLAN. If a RADIUS server subsequently authenticates a new client, but attempts to re-assign the port to a different, untagged VLAN than the one already in use for the previously existing, authenticated client sessions, the connection for the new client will fail.

Example of Untagged VLAN Assignment in a RADIUS-Based Authentication Session

The following example shows how an untagged static VLAN is temporarily assigned to a port for use during an 802.1X authentication session. In the example, an 802.1X-aware client on port A2 has been authenticated by a RADIUS server for access to VLAN 22. However, port A2 is not configured as a member of VLAN 22 but as a member of untagged VLAN 33 as shown in Figure 10-19.

For example, suppose that a RADIUS-authenticated, 802.1X-aware client on port A2 requires access to VLAN 22, but VLAN 22 is configured for no access on port A2, and VLAN 33 is configured as untagged on port A2:

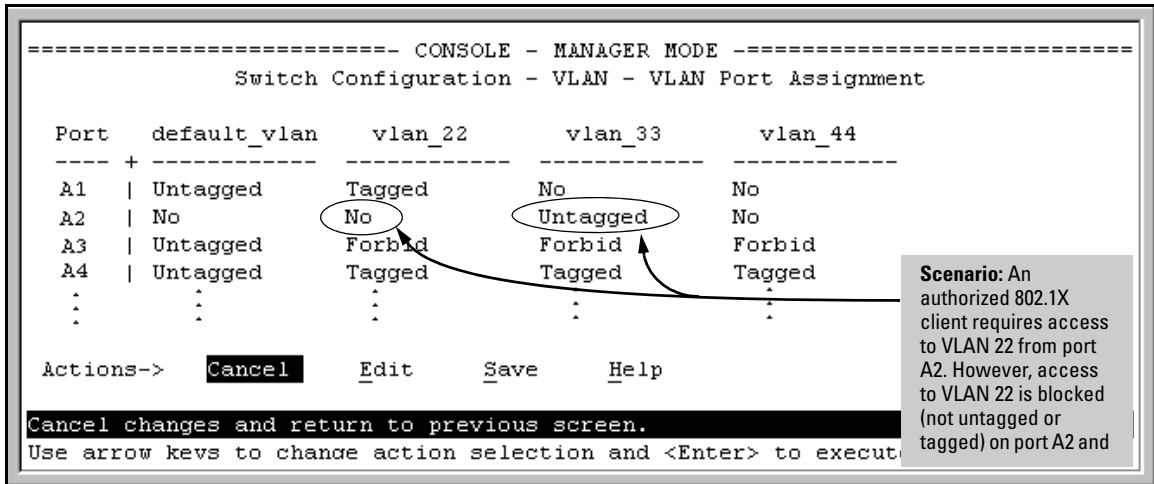


Figure 10-19. Example of an Active VLAN Configuration

In Figure 10-19, if RADIUS authorizes an 802.1X client on port A2 with the requirement that the client use VLAN 22, then:

- VLAN 22 becomes available as Untagged on port A2 for the duration of the session.
- VLAN 33 becomes unavailable to port A2 for the duration of the session (because there can be only one untagged VLAN on any port).

To view the temporary VLAN assignment as a change in the active configuration, use the **show vlan <vlan-id>** command as shown in Figure 10-18 where **<vlan-id>** is the (static or dynamic) VLAN used in the authenticated client session.

Configuring Port-Based and User-Based Access Control (802.1X)

How RADIUS/802.1X Authentication Affects VLAN Operation

```
ProCurve(config)# show vlan 22
Status and Counters - VLAN Information - Ports - VLAN 22
802.1Q VLAN ID : 22
Name           : vlan 22
Status        : Static

Port Information Mode      Unknown VLAN Status
-----
A1              Tagged     Learn          Up
A2              (802.1X)   Learn          Up
A4              Tagged     Learn          Up
.               .          .
.               .          .
.               .          .

Overridden Port VLAN configuration
Port  Mode
----  ---
A2    No
```

This entry shows that port A2 is temporarily untagged on VLAN 22 for an 802.1X session. This is to accommodate an 802.1X client's access, authenticated by a RADIUS server, where the server included an instruction to put the client's access on VLAN 22.

Note: With the current VLAN configuration (figure 10-19), the only time port A2 appears in this **show vlan 22** listing is during an 802.1X session with an attached client. Otherwise, port A2 is not listed.

Figure 10-20.The Active Configuration for VLAN 22 Temporarily Changes for the 802.1X Session

However, as shown in Figure 10-19, because VLAN 33 is configured as untagged on port A2 and because a port can be untagged on only one VLAN, port A2 loses access to VLAN 33 for the duration of the 802.1X session on VLAN 22.

You can verify the temporary loss of access to VLAN 33 by entering the **show vlan 33** command as shown in Figure 10-21.

Even though port A2 is configured as Untagged on (static) VLAN 33 (see figure 10-19), it does not appear in the VLAN 33 listing while the 802.1X session is using VLAN 22 in the Untagged status. However, after the 802.1X session with VLAN 22 ends, the active configuration returns port A2 to VLAN 33.

```
ProCurve# show vlan 33
Status and Counters - VLAN Information - Ports - VLAN 33
802.1Q VLAN ID : 33
Name           : VLAN_33
Status        : Static

Port Information Mode      Unknown VLAN Status
-----
A4              Tagged     Learn          Up

Overridden Port VLAN configuration
Port  Mode
----  ---
A2    Untagged
```

Figure 10-21.The Active Configuration for VLAN 33 Temporarily Drops Port 22 for the 802.1X Session

When the 802.1X client's session on port A2 ends, the port removes the temporary untagged VLAN membership. The static VLAN (VLAN 33) that is "permanently" configured as untagged on the port becomes available again. Therefore, when the RADIUS-authenticated 802.1X session on port A2 ends, VLAN 22 access on port A2 also ends, and the untagged VLAN 33 access on port A2 is restored as shown in Figure 10-22.

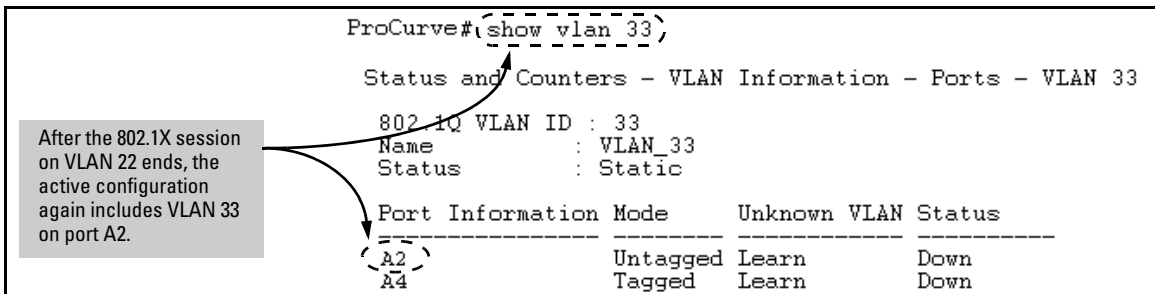


Figure 10-22. The Active Configuration for VLAN 33 Restores Port A2 After the 802.1X Session Ends

Enabling the Use of GVRP-Learned Dynamic VLANs in Authentication Sessions

Syntax: aaa port-access gvrp-vlans

Enables the use of dynamic VLANs (learned through GVRP) in the temporary untagged VLAN assigned by a RADIUS server on an authenticated port in an 802.1X, MAC, or Web authentication session.

*Enter the **no** form of this command to disable the use of GVRP-learned VLANs in an authentication session.*

For information on how to enable a switch to dynamically create 802.1Q-compliant VLANs, see the chapter on "GVRP" in the Advanced Traffic Management Guide.

Notes:

1. If a port is assigned as a member of an untagged dynamic VLAN, the dynamic VLAN configuration must exist at the time of authentication and GVRP for port-access authentication must be enabled on the switch.

If the dynamic VLAN does not exist or if you have not enabled the use of a dynamic VLAN for authentication sessions on the switch, the authentication fails.

Syntax: aaa port-access gvrp-vlans
—Continued—

*2. After you enable dynamic VLAN assignment in an authentication session, it is recommended that you use the **interface unknown-vlans** command on a per-port basis to prevent denial-of-service attacks. The **interface unknown-vlans** command allows you to:*

- *Disable the port from sending advertisements of existing GVRP-created VLANs on the switch.*
- *Drop all GVRP advertisements received on the port.*

For more information, refer to the chapter on “GVRP” in the Advanced Traffic Management Guide.

*3. If you disable the use of dynamic VLANs in an authentication session using the **no aaa port-access gvrp-vlans** command, client sessions that were authenticated with a dynamic VLAN continue and are not deauthenticated.*

(This behavior differs from how static VLAN assignment is handled in an authentication session. If you remove the configuration of the static VLAN used to create a temporary client session, the 802.1X, MAC, or Web authenticated client is deauthenticated.)

However, if a RADIUS-configured dynamic VLAN used for an authentication session is deleted from the switch through normal GVRP operation (for example, if no GVRP advertisements for the VLAN are received on any switch port), authenticated clients using this VLAN are deauthenticated.

Note

Any port VLAN-ID changes you make on 802.1X-aware ports during an 802.1X-authenticated session do not take effect until the session ends.

With GVRP enabled, a temporary, untagged static VLAN assignment created on a port by 802.1X authentication is advertised as an existing VLAN. If this temporary VLAN assignment causes the switch to disable a configured (untagged) static VLAN assignment on the port, then the disabled VLAN assignment is not advertised. When the 802.1X session ends, the switch:

- Eliminates and ceases to advertise the temporary VLAN assignment.
 - Re-activates and resumes advertising the temporarily disabled VLAN assignment.
-

Messages Related to 802.1X Operation

Table 10-4. 802.1X Operating Messages

Message	Meaning
Port <port-list> is not an authenticator.	The ports in the port list have not been enabled as 802.1X authenticators. Use this command to enable the ports as authenticators: <pre>ProCurve(config)# aaa port-access authenticator e 10</pre>
Port <port-list> is not a supplicant.	Occurs when there is an attempt to change the supplicant configuration on a port that is not currently enabled as a supplicant. Enable the port as a supplicant and then make the desired supplicant configuration changes. Refer to “Enabling a Switch Port as a Supplicant” on page 10-49.
No server(s) responding.	This message can appear if you configured the switch for EAP-RADIUS or CHAP-RADIUS authentication, but the switch does not receive a response from a RADIUS server. Ensure that the switch is configured to access at least one RADIUS server. (Use show radius .) If you also see the message <code>Can't reach RADIUS server <x.x.x.x></code> , try the suggestions listed for that message (page 5-53).
LACP has been disabled on 802.1X port(s).	To maintain security, LACP is not allowed on ports configured for 802.1X authenticator operation. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port(s), and enables 802.1X on that port.
Error configuring port <port-number>: LACP and 802.1X cannot be run together.	Also, the switch will not allow you to configure LACP on a port on which port access (802.1X) is enabled.

Configuring and Monitoring Port Security

Contents

Overview	11-3
Port Security	11-4
Basic Operation	11-4
Eavesdrop Protection	11-5
Blocking Unauthorized Traffic	11-5
Trunk Group Exclusion	11-6
Planning Port Security	11-7
Port Security Command Options and Operation	11-8
Port Security Display Options	11-8
Configuring Port Security	11-12
Retention of Static Addresses	11-17
MAC Lockdown	11-22
Differences Between MAC Lockdown and Port Security	11-24
MAC Lockdown Operating Notes	11-25
Deploying MAC Lockdown	11-26
MAC Lockout	11-26
Port Security and MAC Lockout	11-29
Web: Displaying and Configuring Port Security Features	11-30
Reading Intrusion Alerts and Resetting Alert Flags	11-30
Notice of Security Violations	11-30
How the Intrusion Log Operates	11-31
Keeping the Intrusion Log Current by Resetting Alert Flags	11-32
Menu: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags	11-33
CLI: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags	11-34
Using the Event Log To Find Intrusion Alerts	11-36

Web: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags	11-37
Operating Notes for Port Security	11-38

Overview

Feature	Default	Menu	CLI	Web
Displaying Current Port Security	n/a	—	page 11-8	page 11-30
Configuring Port Security	disabled	—	page 11-12	page 11-30
Retention of Static Addresses	n/a	—	page 11-17	n/a
MAC Lockdown	disabled	—	page 11-22	
MAC Lockout	disabled	—	page 11-26	
Intrusion Alerts and Alert Flags	n/a	page 11-36	page 11-34	page 11-37

Port Security (Page 11-4). This feature enables you to configure each switch port with a unique list of the MAC addresses of devices that are authorized to access the network through that port. This enables individual ports to detect, prevent, and log attempts by unauthorized devices to communicate through the switch.

Note

This feature does not prevent intruders from receiving broadcast and multi-cast traffic. Also, Port Security and MAC Lockdown are mutually exclusive on a switch. If one is enabled, then the other cannot be used.

MAC Lockdown (Page 11-22). This feature, also known as “Static Addressing”, is used to prevent station movement and MAC address “hijacking” by allowing a given MAC address to use only an assigned port on the switch. MAC Lockdown also restricts the client device to a specific VLAN. (See also the **Note**, above.)

MAC Lockout (Page 11-26). This feature enables you to block a specific MAC address so that the switch drops all traffic to or from the specified address.

Port Security

Basic Operation

Default Port Security Operation. The default port security setting for each port is off, or “continuous”. That is, any device can access a port without causing a security reaction.

Intruder Protection. A port that detects an “intruder” blocks the intruding device from transmitting to the network through that port.

Eavesdrop Protection. Using either the port-security command or the switch’s web browser interface to enable port security on a given port automatically enables eavesdrop prevention on that port.

General Operation for Port Security. On a per-port basis, you can configure security measures to block unauthorized devices, and to send notice of security violations. Once port security is configured, you can then monitor the network for security violations through one or more of the following:

- Alert flags that are captured by network management tools such as ProCurve Manager (PCM and PCM+)
- Alert Log entries in the switch’s web browser interface
- Event Log entries in the console interface
- Intrusion Log entries in the menu interface, CLI, or web browser interface

For any port, you can configure the following:

- **Action:** Used when a port detects an intruder. Specifies whether to send an SNMP trap to a network management station and whether to disable the port.
- **Address Limit:** Sets the number of authorized MAC addresses allowed on the port.
- **Learn-Mode:** Specify how the port acquires authorized addresses.
 - **Continuous:** Allows the port to learn addresses from inbound traffic from any connected device. This is the default setting.
 - **Limited-Continuous:** Sets a finite limit (1 - 32) to the number of learned addresses allowed per port.

- **Static:** Enables you to set a fixed limit on the number of MAC addresses authorized for the port and to specify some or all of the authorized addresses. (If you specify only some of the authorized addresses, the port learns the remaining authorized addresses from the traffic it receives from connected devices.)
- **Configured:** Requires that you specify all MAC addresses authorized for the port. The port is not allowed to learn addresses from inbound traffic.
- **Authorized (MAC) Addresses:** Specify up to eight devices (MAC addresses) that are allowed to send inbound traffic through the port. This feature:
 - Closes the port to inbound traffic from any unauthorized devices that are connected to the port.
 - Provides the option for sending an SNMP trap notifying of an attempted security violation to a network management station and, optionally, disables the port. (For more on configuring the switch for SNMP management, see “Trap Receivers and Authentication Traps” in the *Management and Configuration Guide* for your switch.)
- **Port Access:** Allows only the MAC address of a device authenticated through the switch’s 802.1X Port-Based access control. Refer to chapter 10, Configuring Port-Based and User-Based Access Control (802.1X).

For configuration details, refer to “Configuring Port Security” on page 11-12.

Eavesdrop Protection

Configuring port security on a given switch port automatically enables eavesdrop protection for that port. This prevents use of the port to flood unicast packets addressed to MAC addresses unknown to the switch. This blocks unauthorized users from eavesdropping on traffic intended for addresses that have aged-out of the switch’s address table. (Eavesdrop prevention does not affect multicast and broadcast traffic, meaning that the switch floods these two traffic types out a given port regardless of whether port security is enabled on that port.)

Blocking Unauthorized Traffic

Unless you configure the switch to disable a port on which a security violation is detected, the switch security measures block unauthorized traffic without disabling the port. This implementation enables you to apply the security

configuration to ports on which hubs, switches, or other devices are connected, and to maintain security while also maintaining network access to authorized users. For example:

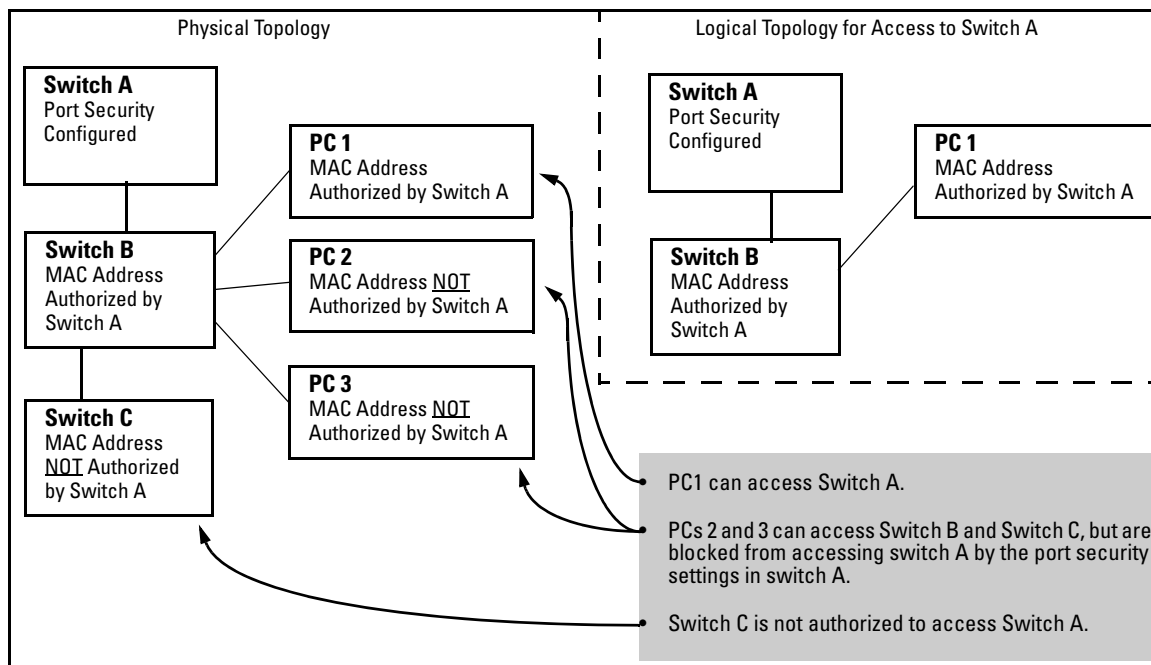


Figure 11-1. Example of How Port Security Controls Access

Note

Broadcast and Multicast traffic is always allowed, and can be read by intruders connected to a port on which you have configured port security.

Trunk Group Exclusion

Port security does not operate on either a static or dynamic trunk group. If you configure port security on one or more ports that are later added to a trunk group, the switch will reset the port security parameters for those ports to the factory-default configuration. (Ports configured for either Active or Passive LACP, and which are not members of a trunk, can be configured for port security.)

Planning Port Security

1. Plan your port security configuration and monitoring according to the following:
 - a. On which ports do you want port security?
 - b. Which devices (MAC addresses) are authorized on each port?
 - c. For each port, what security actions do you want? (The switch automatically blocks intruders detected on that port from transmitting to the network.) You can configure the switch to (1) send intrusion alarms to an SNMP management station and to (2) optionally disable the port on which the intrusion was detected.
 - d. How do you want to learn of the security violation attempts the switch detects? You can use one or more of these methods:
 - Through network management (That is, do you want an SNMP trap sent to a net management station when a port detects a security violation attempt?)
 - Through the switch's Intrusion Log, available through the CLI, menu, and web browser interface
 - Through the Event Log (in the menu interface or through the CLI **show log** command)
2. Use the CLI or web browser interface to configure port security operating and address controls. The following table describes the parameters.

Port Security Command Options and Operation

Port Security Commands Used in This Section

show port-security	11-9
show mac-address	
port-security	11-12
< <i>port-list</i> >	11-12
learn-mode	11-12
address-limit	11-15
mac-address	11-16
action	11-16
clear-intrusion-flag	11-17
no port-security	11-17

This section describes the CLI port security command and how the switch acquires and maintains authorized addresses.

Note

Use the global configuration level to execute port-security configuration commands.

Port Security Display Options

You can use the CLI to display the current port-security settings and to list the currently authorized MAC addresses the switch detects on one or more ports.

Displaying Port Security Settings.

Syntax: show port-security
show port-security <port number>
show port-security [<port number>-<port number>]. . [<port number>]

The CLI uses the same command to provide two types of port security listings:

- All ports on the switch with their Learn Mode and (alarm) Action
- Only the specified ports with their Learn Mode, Address Limit, (alarm) Action, and Authorized Addresses

*Without port parameters, **show port-security** displays Operating Control settings for all ports on a switch.*

ProCurve(config)# show port-security				
Port Security				
Port		Learn Mode		Action
----		-----	+	-----
A1	1	Static		Send Alarm, Disable Port
A2	2	Static		Send Alarm, Disable Port
A3	3	Static		Send Alarm
A4	4	Static		Send Alarm
A5	5	Static		Send Alarm
A6	6	Static		Send Alarm
A7	7	Continuous		None
A8	8	Continuous		None

Figure 11-2. Example Port Security Listing (Ports A7 and A8 Show the Default Setting)

With port numbers included in the command, **show port-security** displays Learn Mode, Address Limit, (alarm) Action, and Authorized Addresses for the specified ports on a switch. The following example lists the full port security configuration for a single port:

```
ProCurve(config)# show port-security A3
Port Security
  Port : A3
  Learn Mode : Static           Address Limit : 1
  Action : Send Alarm

  Authorized Addresses
  -----
  00906d-fdcc00
```

Figure 11-3. Example of the Port Security Configuration Display for a Single Port

The next example shows the option for entering a range of ports, including a series of non-contiguous ports. Note that no spaces are allowed in the port number portion of the command string:

```
ProCurve(config)# show port-security A1-A3,A6,A8
```

Listing Authorized and Detected MAC Addresses.

Syntax: show mac-address [port-list | mac-address | vlan < vid >]

*Without an optional parameter, **show mac-address** lists the authorized MAC addresses that the switch detects on all ports.*

mac-address: *Lists the specified MAC address with the port on which it is detected as an authorized address.*

port list: *Lists the authorized MAC addresses detected on the specified port(s).*

vlan < vid >: *Lists the authorized MAC addresses detected on ports belonging to the specified VLAN.*

```
ProCurve(config)# show mac-address

Status and Counters - Port Address Table

MAC Address    Located on Port
-----
0004ea-84d980  7
0004ea-84d9ee  7
000a57-4d8d40  5
:
:
00a0c9-f1786f  5

ProCurve(config)# show mac-address 7

Status and Counters - Port Address Table - 7

MAC Address
-----
0004ea-84d980
0004ea-84d9ee

ProCurve(config)# show mac-address 000a57-4d8d40

Status and Counters - Address Table - 000a57-4d8d40

MAC Address : 000a57-4d8d40
Located on Port : 5

ProCurve(config)# show mac-address vlan 1

Status and Counters - Address Table - VLAN 1

MAC Address    Located on Port
-----
0004ea-84d980  7
0004ea-84d9ee  7
000a57-4d8d40  5
:
:
00a0c9-f1786f  5
```

Figure 11-4. Examples of Show Mac-Address Outputs

Configuring Port Security

Using the CLI, you can:

- Configure port security and edit security settings.
- Add or delete devices from the list of authorized addresses for one or more ports.
- Clear the Intrusion flag on specific ports

Syntax: port-security

[e] <port-list>< learn-mode | address-limit | mac-address | action |
clear-intrusion-flag >

< port-list >: Specifies a list of one or more ports to which the port-security command applies.

learn-mode < continuous | static | port-access | configured | limited-continuous >

For the specified port:

- Identifies the method for acquiring authorized addresses.
- On switches covered in this guide, automatically invokes eavesdrop protection. (Refer to “Eavesdrop Protection” on page 11-5.)

continuous (Default): Appears in the factory-default setting or when you execute **no port-security**. Allows the port to learn addresses from the device(s) to which it is connected. In this state, the port accepts traffic from any device(s) to which it is connected. Addresses learned in the learn continuous mode will “age out” and be automatically deleted if they are not used regularly. The default age time is five minutes.

Addresses learned this way appear in the switch and port address tables and age out according to the **MAC Age Interval** in the System Information configuration screen of the Menu interface or the **show system information** listing. You can set the MAC age out time using the CLI, SNMP, Web, or menu interfaces. For more information on the **mac-age-time** command refer to the chapter titled “Interface Access and System Information” in the Management and Configuration Guide for your switch.

— Continued —

Syntax: port-security (*Continued*)

learn-mode < continuous | static | port-access | configured | limited-continuous > (*Continued*)

static: *Enables you to use the **mac-address** parameter to specify the MAC addresses of the devices authorized for a port, and the **address-limit** parameter (explained below) to specify the number of MAC addresses authorized for the port. You can authorize specific devices for the port, while still allowing the port to accept other, non-specified devices until the device limit has been reached. That is, if you enter fewer MAC addresses than you authorized, the port authorizes the remaining addresses in the order in which it automatically learns them.*

*For example, if you use address-limit to specify three authorized devices, but use **mac-address** to specify only one authorized MAC address, the port adds the one specifically authorized MAC address to its authorized-devices list and the first two additional MAC addresses it detects.*

If, for example:

*You use **mac-address** to authorize MAC address 0060b0-880a80 for port A4.*

*You use **address-limit** to allow three devices on port A4 and the port detects these MAC addresses:*

- 1. 080090-1362f2*
- 2. 00f031-423fc1*
- 3. 080071-0c45a1*
- 4. 0060b0-880a80 (the address you authorized with the **mac-address** parameter)*

In this example port A4 would assume the following list of authorized addresses:

080090-1362f2 (the first address the port detected)

00f031-423fc1 (the second address the port detected)

*0060b0-880a80 (the address you authorized with the **mac-address** parameter)*

*The remaining MAC address detected by the port, 080071-0c45a1, is not allowed and is handled as an intruder. Learned addresses that become authorized do **not** age-out. See also “Retention of Static Addresses” on page 11-17.*

— Continued —

Syntax: port-security (*Continued*)

learn-mode < continuous | static | port-access | configured | limited-continuous > (*Continued*)

Caution: Using the **static** parameter with a device limit greater than the number of MAC addresses specified with **mac-address** can allow an unwanted device to become “authorized”. This is because the port, to fulfill the number of devices allowed by the **address-limit** parameter (see below), automatically adds devices it detects until it reaches the specified limit.

Note: If 802.1X port-access is configured on a given port, then port-security learn-mode must be set to either **continuous** (the default) or **port-access**.

port-access: Enables you to use Port Security with (802.1X) Port-Based Access Control. Refer to chapter 10, *Configuring Port-Based and User-Based Access Control (802.1X)*.

configured: Must specify which MAC addresses are allowed for this port. Range is 1 (default) to 8 and addresses are not ageable. Addresses are saved across reboots.

limited-continuous: Also known as MAC Secure, or “limited” mode. The limited parameter sets a finite limit to the number of learned addresses allowed per port. (You can set the range from 1, the default, to a maximum of 32 MAC addresses which may be learned by each port.)

All addresses are ageable, meaning they are automatically removed from the authorized address list for that port after a certain amount of time. Limited mode and the address limit are saved across reboots, but addresses which had been learned are lost during the reboot process.

Addresses learned in the limited mode are normal addresses learned from the network until the limit is reached, but they are not configurable. (You cannot enter or remove these addresses manually if you are using **learn-mode** with the **limited-continuous** option.)

—Continued—

Syntax: port-security (*Continued*)

*Addresses learned this way appear in the switch and port address tables and age out according to the **MAC Age Interval** in the System Information configuration screen of the Menu interface or the **show system information** listing. You can set the MAC age out time using the CLI, SNMP, Web, or menu interfaces. For more on the **mac-age-time** command, refer to the chapter titled “Interface Access and System Information” in the Management and Configuration Guide for your switch. To set the learn-mode to **limited** use this command syntax:*

```
port-security <port-list> learn-mode limited address-limit  
< 1..32 > action < none | send-alarm | send-disable >
```

*The default address-limit is **1** but may be set for each port to learn up to 32 addresses. The default action is **none**. To see the list of learned addresses for a port use the command:*

```
show mac < port-list >
```

```
address-limit < integer >
```

*When **learn-mode** is set to **static**, **configured**, or **limited-continuous**, the **address-limit** parameter specifies how many authorized devices (MAC addresses) to allow. Range: 1 (the default) to 8 for static and configured modes. For **learn-mode** with the **limited-continuous** option, the range is 1-32 addresses.*

—Continued—

Syntax: port-security (Continued)

mac-address [<mac-addr>] [<mac-addr>] ... [<mac-addr>]

Available for **learn-mode** with the, **static**, **configured**, or **limited-continuous** option. Allows up to eight authorized devices (MAC addresses) per port, depending on the value specified in the **address-limit** parameter. The **mac-address limited-continuous** mode allows up to 32 authorized MAC addresses per port.

If you use **mac-address** with **static**, but enter fewer devices than you specified in the **address-limit** field, the port accepts not only your specified devices, but also as many other devices as it takes to reach the device limit. For example, if you specify four devices, but enter only two MAC addresses, the port will accept the first two non-specified devices it detects, along with the two specifically authorized devices. Learned addresses that become authorized do **not** age-out. See also “Retention of Static Addresses” on page 11-17.

action < none | send-alarm | send-disable >

Specifies whether an SNMP trap is sent to a network management station when Learn Mode is set to **static** and the port detects an unauthorized device, or when Learn Mode is set to continuous and there is an address change on a port.

none: Prevents an SNMP trap from being sent. **none** is the default value.

send-alarm: Sends an intrusion alarm. Causes the switch to send an SNMP trap to a network management station.

send-disable: Sends alarm and disables the port. Available only in the **static**, **port-access**, **configured**, or **limited learn** modes. Causes the switch to send an SNMP trap to a network management station and disable the port. If you subsequently re-enable the port without clearing the port's intrusion flag, the port blocks further intruders, but the switch will not disable the port again until you reset the intrusion flag. See the Note on 11-32.

For information on configuring the switch for SNMP management, refer to the Management and Configuration Guide for your switch.

—Continued—

Syntax: port-security (*Continued*)

clear-intrusion-flag

Clears the intrusion flag for a specific port. (See “Reading Intrusion Alerts and Resetting Alert Flags” on page 11-30.)

no port-security <port-list> mac-address <mac-addr> [<mac-addr>
<mac-addr>]

Removes the specified learned MAC address(es) from the specified port.

Retention of Static Addresses

Static MAC addresses do not age-out. MAC addresses learned by using **learn-mode continuous** or **learn-mode limited-continuous** age out according to the currently configured MAC age time. (For information on the **mac-age-time** command, refer to the chapter titled “Interface Access and System Information” in the *Management and Configuration Guide* for your switch.

Learned Addresses. In the following two cases, a port in Static learn mode retains a learned MAC address even if you later reboot the switch or disable port security for that port:

- The port learns a MAC address after you configure the port for Static learn mode in both the startup-config file and the running-config file (by executing the **write memory** command).
- The port learns a MAC address after you configure the port for Static learn mode in only the running-config file and, after the address is learned, you execute **write memory** to configure the startup-config file to match the running-config file.

To remove an address learned using either of the preceding methods, do one of the following:

- Delete the address by using **no port-security < port-number > mac-address < mac-addr >**.
- Download a configuration file that does not include the unwanted MAC address assignment.
- Reset the switch to its factory-default configuration.

Assigned/Authorized Addresses. : If you manually assign a MAC address (using **port-security <port-number> address-list <mac-addr>**) and then execute **write memory**, the assigned MAC address remains in memory until you do one of the following:

- Delete it by using **no port-security <port-number> mac-address <mac-addr>**.
- Download a configuration file that does not include the unwanted MAC address assignment.
- Reset the switch to its factory-default configuration.

Specifying Authorized Devices and Intrusion Responses. This example configures port A1 to automatically accept the first device (MAC address) it detects as the only authorized device for that port. (The default device limit is 1.) It also configures the port to send an alarm to a network management station and disable itself if an intruder is detected on the port.

```
ProCurve(config)# port-security a1 learn-mode static  
action send-disable
```

The next example does the same as the preceding example, except that it specifies a MAC address of 0c0090-123456 as the authorized device instead of allowing the port to automatically assign the first device it detects as an authorized device.

```
ProCurve(config)# port-security a1 learn-mode static  
mac-address 0c0090-123456 action send-disable
```

This example configures port A5 to:

- Allow two MAC addresses, 00c100-7fec00 and 0060b0-889e00, as the authorized devices.
- Send an alarm to a management station if an intruder is detected on the port, but allow the intruder access to the network.

```
ProCurve(config)# port-security a5 learn-mode static  
address-limit 2 mac-address 00c100-7fec00 0060b0-889e00  
action send-alarm
```

If you manually configure authorized devices (MAC addresses) and/or an alarm action on a port, those settings remain unless you either manually change them or the switch is reset to its factory-default configuration. You can “turn off” authorized devices on a port by configuring the port to continuous Learn Mode, but subsequently reconfiguring the port to static Learn Mode restores those authorized devices.

Adding an Authorized Device to a Port. To simply add a device (MAC address) to a port's existing Authorized Addresses list, enter the port number with the **mac-address** parameter and the device's MAC address. *This assumes that Learn Mode is set to **static** and the Authorized Addresses list is not full* (as determined by the current Address Limit value). For example, suppose port A1 allows two authorized devices, but has only one device in its Authorized Address list:

```
ProCurve(config)# show port-security a1

Port Security
  Port : A1
  Learn Mode [Continuous] : Static      Address Limit [1] : 2
  Action [None] : None

  Authorized Addresses
  -----
  0c0090-123456
```

Although the Address Limit is set to 2, only one device has been authorized for this port. In this case you can add another without having to also increase the Address Limit.

The Address Limit has not been reached.

Figure 11-5. Example of Adding an Authorized Device to a Port

With the above configuration for port A1, the following command adds the 0c0090-456456 MAC address as the second authorized address.

```
ProCurve(config)# port-security a1 mac-address 0c0090-456456
```

After executing the above command, the security configuration for port A1 would be:

```
ProCurve(config)# show port-security a1

Port Security
  Port : A1
  Learn Mode [Continuous] : Static      Address Limit [1] : 2
  Action [None] : None

  Authorized Addresses
  -----
  0c0090-123456
  0c0090-456456
```

The Address Limit has been reached.

Figure 11-6. Example of Adding a Second Authorized Device to a Port

(The message **Inconsistent value** appears if the new MAC address exceeds the current Address Limit or specifies a device that is already on the list. Note that if you change a port from static to continuous learn mode, the port retains in memory any authorized addresses it had while in static mode. If you subsequently attempt to convert the port back to static mode with the same authorized address(es), the `Inconsistent value` message appears because the port already has the address(es) in its “Authorized” list.)

If you are adding a device (MAC address) to a port on which the Authorized Addresses list is already full (as controlled by the port’s current Address Limit setting), then you must increase the Address Limit in order to add the device, even if you want to replace one device with another. Using the CLI, you can simultaneously increase the limit and add the MAC address with a single command. For example, suppose port A1 allows one authorized device and already has a device listed:

```
HPswitch(config)# show port-security a1
Port Security
  Port : A1
  Learn Mode [Continuous] : Static    Address Limit [1]:1
  Action [None] : None

  Authorized Addresses
  -----
  0c0090-123456
```

Figure 11-7. Example of Port Security on Port A1 with an Address Limit of “1”

To add a second authorized device to port A1, execute a **port-security** command for port A1 that raises the address limit to 2 and specifies the additional device’s MAC address. For example:

```
ProCurve(config)# port-security a1 mac-address 0c0090-456456 address-limit 2
```

Removing a Device From the “Authorized” List for a Port. This command option removes unwanted devices (MAC addresses) from the Authorized Addresses list. (An Authorized Address list is available for each port for which Learn Mode is currently set to “Static”. Refer to the command syntax listing under “Configuring Port Security” on page 11-12.)

Caution

When learn mode is set to static, the Address Limit (address-limit) parameter controls how many devices are allowed in the Authorized Addresses (**mac-address**) for a given port. If you remove a MAC address from the Authorized Addresses list without also reducing the Address Limit by 1, the port may subsequently detect and accept as authorized a MAC address that you do not intend to include in your Authorized Address list. Thus, if you use the CLI to remove a device that is no longer authorized, it is recommended that you first reduce the Address Limit (**address-limit**) integer by 1, as shown below. This prevents the possibility of the same device or another unauthorized device on the network from automatically being accepted as “authorized” for that port.

To remove a device (MAC address) from the “Authorized” list and when the current number of devices equals the Address Limit value, you should first reduce the Address Limit value by 1, then remove the unwanted device.

Note

You can reduce the address limit below the number of currently authorized addresses on a port. This enables you to subsequently remove a device from the “Authorized” list without opening the possibility for an unwanted device to automatically become authorized.

For example, suppose port A1 is configured as shown below and you want to remove 0c0090-123456 from the Authorized Address list:

```
ProCurve(config)# show port-security a1
Port Security
Port : A1
Learn Mode [Continuous] : Static      Address Limit [1] : 2
Action [None] : None

Authorized Addresses
-----
0c0090-123456
0c0090-456456
```

When removing 0c0090-123456, first reduce the Address Limit by 1 to prevent the port from automatically adding another device that it detects on the network.

Figure 11-8. Example of Two Authorized Addresses on Port A1

The following command serves this purpose by removing 0c0090-123456 and reducing the Address Limit to 1:

```
ProCurve(config)# port-security a1 address-limit 1
ProCurve(config)# no port-security a1 mac-address 0c0090-123456
```

The above command sequence results in the following configuration for port A1:

```
ProCurve(config)# show port-sec a1
Port Security
  Port : A1
  Learn Mode : Static           Address Limit : 1
  Action : None
  Authorized Addresses
  -----
  0c0090-456456
```

Figure 11-9. Example of Port A1 After Removing One MAC Address

MAC Lockdown

MAC Lockdown, also known as “static addressing,” is the permanent assignment of a given MAC address (and VLAN, or Virtual Local Area Network) to a specific port on the switch. MAC Lockdown is used to prevent station movement and MAC address hijacking. It also controls address learning on the switch. When configured, the MAC Address can only be used on the assigned port and the client device will only be allowed on the assigned VLAN.

Note

Port security and MAC Lockdown are mutually exclusive on a given port. You can either use port security *or* MAC Lockdown, but never both at the same time on the same port.

Syntax: [no] static-mac < mac-addr > vlan < vid > interface < port-number >

You will need to enter a separate command for each MAC/VLAN pair you wish to lock down. If you do not specify a VLAN ID (VID) the switch inserts a VID of “1”.

How It Works. When a device’s MAC address is locked down to a port (typically in a pair with a VLAN) all information sent to that MAC address must go through the locked-down port. If the device is moved to another port it cannot receive data. Traffic to the designated MAC address goes only to the allowed port, whether the device is connected to it or not.

MAC Lockdown is useful for preventing an intruder from “hijacking” a MAC address from a known user in order to steal data. Without MAC Lockdown, this will cause the switch to learn the address on the malicious user’s port, allowing the intruder to steal the traffic meant for the legitimate user.

MAC Lockdown ensures that traffic intended for a specific MAC address can only go through the one port which is supposed to be connected to that MAC address. It does not prevent intruders from transmitting packets with the locked MAC address, but it does prevent responses to those packets from going anywhere other than the locked-down port. Thus TCP connections cannot be established. Traffic sent to the locked address cannot be hijacked and directed out the port of the intruder.

If the device (computer, PDA, wireless device) is moved to a different port on the switch (by reconnecting the Ethernet cable or by moving the device to an area using a wireless access point connected to a different port on that same switch), the port will detect that the MAC Address is not on the appropriate port and will continue to send traffic out the port to which the address was locked.

Once a MAC address is configured for one port, you cannot perform port security using the same MAC address on any other port on that same switch.

You cannot lock down a single MAC Address/VLAN pair to more than one port; however you can lock down multiple different MAC Addresses to a single port on the same switch.

Stations can move from the port to which their MAC address is locked to other parts of the network. They can send, but will not receive data if that data must go through the locked down switch. Please note that if the device moves to a distant part of the network where data sent to its MAC address never goes through the locked down switch, it may be possible for the device to have full two-way communication. For full and complete lockdown network-wide all switches must be configured appropriately.

Other Useful Information. Once you lock down a MAC address/VLAN pair on one port that pair cannot be locked down on a different port.

You cannot perform MAC Lockdown and 802.1X authentication on the same port or on the same MAC address. MAC Lockdown and 802.1X authentication are mutually exclusive.

Lockdown is permitted on static trunks (manually configured link aggregations).

Differences Between MAC Lockdown and Port Security

Because port-security relies upon MAC addresses, it is often confused with the MAC Lockdown feature. However, MAC Lockdown is a completely different feature and is implemented on a different architecture level.

Port security maintains a list of allowed MAC addresses on a per-port basis. An address can exist on multiple ports of a switch. Port security deals with MAC addresses only while MAC Lockdown specifies both a MAC address and a VLAN for lockdown.

MAC Lockdown, on the other hand, is not a “list.” It is a global parameter on the switch that takes precedence over any other security mechanism. The MAC Address will only be allowed to communicate using one specific port on the switch.

MAC Lockdown is a good replacement for port security to create tighter control over MAC addresses and which ports they are allowed to use (only one port per MAC Address on the same switch in the case of MAC Lockdown). (You can still use the port for other MAC addresses, but you cannot use the locked down MAC address on other ports.)

Using only port security the MAC Address could still be used on another port on the same switch. MAC Lockdown, on the other hand, is a clear one-to-one relationship between the MAC Address and the port. Once a MAC address has been locked down to a port it cannot be used on another port on the same switch.

The switch does not allow MAC Lockdown and port security on the same port.

MAC Lockdown Operating Notes

Limits. There is a limit of 500 MAC Lockdowns that you can safely code per switch. To truly lock down a MAC address it would be necessary to use the MAC Lockdown command for every MAC Address and VLAN ID on every switch. In reality few network administrators will go to this length, but it is important to note that just because you have locked down the MAC address and VID for a single switch, the device (or a hacker “spoofing” the MAC address for the device) may still be able to use another switch which hasn’t been locked down.

Event Log Messages. If someone using a locked down MAC address is attempting to communicate using the wrong port the “move attempt” generates messages in the log file like this:

Move attempt (lockdown) logging:

```
W 10/30/03 21:33:43 maclock: module A: Move 0001e6-1f96c0  
to A15 denied
```

```
W 10/30/03 21:33:48 maclock: module A: Move 0001e6-1f96c0  
to A15 denied
```

```
W 10/30/03 21:33:48 maclock: module A: Ceasing move-denied  
logs for 5m
```

These messages in the log file can be useful for troubleshooting problems. If you are trying to connect a device which has been locked down to the wrong port, it will not work but it will generate error messages like this to help you determine the problem.

Limiting the Frequency of Log Messages. The first move attempt (or intrusion) is logged as you see in the example above. Subsequent move attempts send a message to the log file also, but message throttling is imposed on the logging on a per-module basis. What this means is that the logging system checks again after the first 5 minutes to see if another attempt has been made to move to the wrong port. If this is the case the log file registers the most recent attempt and then checks again after one hour. If there are no further attempts in that period then it will continue to check every 5 minutes. If another attempt was made during the one hour period then the log resets itself to check once a day. The purpose of rate-limiting the log messaging is to prevent the log file from becoming too full. You can also configure the switch to send the same messages to a Syslog server. Refer to “Debug and Syslog Messaging Operation” in appendix C of the *Management and Configuration Guide* for your switch.

Deploying MAC Lockdown

When you deploy MAC Lockdown you need to consider how you use it within your network topology to ensure security. In some cases where you are using techniques such as Spanning Tree Protocol (STP) to speed up network performance by providing multiple paths for devices, using MAC Lockdown either will not work or else it defeats the purpose of having multiple data paths.

The purpose of using MAC Lockdown is to prevent a malicious user from “hijacking” an approved MAC address so they can steal data traffic being sent to that address.

As we have seen, MAC Lockdown can help prevent this type of hijacking by making sure that all traffic to a specific MAC address goes only to the proper port on a switch which is supposed to be connected to the real device bearing that MAC address.

However, you can run into trouble if you incorrectly try to deploy MAC Lockdown in a network that uses multiple path technology, like Spanning Tree.

Caution

Using MAC Lockdown still does not protect against a hijacker *within the core*! In order to protect against someone spoofing the MAC Address on a server inside the Core Network, you would have to lock down each and every switch inside the Core Network as well, not just on the edge.

MAC Lockout

MAC Lockout involves configuring a MAC address on all ports and VLANs for a switch so that any traffic to or from the “locked-out” MAC address will be dropped. This means that all data packets addressed to or from the given address are stopped by the switch. MAC Lockout is implemented on a per switch assignment.

You can think of MAC Lockout as a simple blacklist. The MAC address is locked out on the switch and on all VLANs. No data goes out or in from the blacklisted MAC address to a switch using MAC Lockout.

To fully lock out a MAC address from the network it would be necessary to use the MAC Lockout command on all switches.

To use MAC Lockout you must first know the MAC Address you wish to block.

Syntax: [no] `lockout-mac < mac-address >`

How It Works. Let's say a customer knows there are unauthorized wireless clients who should not have access to the network. The network administrator "locks out" the MAC addresses for the wireless clients by using the MAC Lockout command (**`lockout-mac <mac-address>`**). When the wireless clients then attempt to use the network, the switch recognizes the intruding MAC addresses and prevents them from sending or receiving data on that network.

If a particular MAC address can be identified as unwanted on the switch then that MAC Address can be disallowed on all ports on that switch with a single command. You don't have to configure every single port—just perform the command on the switch and it is effective for all ports.

MAC Lockout overrides MAC Lockdown, port security, and 802.1X authentication.

You cannot use MAC Lockout to lock:

- Broadcast or Multicast Addresses (Switches do not learn these)
- Switch Agents (The switch's own MAC Address)

There are limits for the number of VLANs and Lockout MACs that can be configured concurrently as all use MAC table entries. The limits are shown below.

Table 11-10. Limits on Lockout MACs

# VLANs	# Lockout MACs
<= 1024	16
1025-2048	8

If someone using a locked out MAC address tries to send data through the switch a message is generated in the log file:

Lockout logging format:

```
W 10/30/03 21:35:15 maclock: module A: 0001e6-1f96c0
detected on port A15
W 10/30/03 21:35:18 maclock: module A: 0001e6-1f96c0
detected on port A15
W 10/30/03 21:35:18 maclock: module A: Ceasing lock-out
logs for 5m
```

As with MAC Lockdown a rate limiting algorithm is used on the log file so that it does not become overlogged with error messages. (Refer to “Limiting the Frequency of Log Messages” on page 11-25.)

Port Security and MAC Lockout

MAC Lockout is independent of port-security and in fact will override it. MAC Lockout is preferable to port-security to stop access from known devices because it can be configured for all ports on the switch with one command.

It is possible to use MAC Lockout in conjunction with port-security. You can use MAC Lockout to lock out a single address—deny access to a specific device—but still allow the switch some flexibility in learning other MAC Addresses. Be careful if you use both together, however:

- If a MAC Address is locked out and appears in a static learn table in port-security, the apparently “authorized” address will still be locked out anyway.
- MAC entry configurations set by port security will be kept even if MAC Lockout is configured and the original port security settings will be honored once the Lockout is removed.
- A port security static address is permitted to be a lockout address. In that case (MAC Lockout), the address will be locked out (SA/DA drop) even though it’s an “authorized” address from the perspective of port security.
- When MAC Lockout entries are deleted, port security will then re-learn the address as needed later on.

Web: Displaying and Configuring Port Security Features

1. Click on the **Security** tab.
2. Click on **[Port Security]**.
3. Select the settings you want and, if you are using the Static Learn Mode, add or edit the Authorized Addresses field.
4. Implement your new data by clicking on **[Apply Changes]**.

To access the web-based Help provided for the switch, click on **[?]** in the web browser screen.

Reading Intrusion Alerts and Resetting Alert Flags

Notice of Security Violations

When the switch detects an intrusion on a port, it sets an “alert flag” for that port and makes the intrusion information available as described below. *While the switch can detect additional intrusions for the same port, it does not list the next chronological intrusion for that port in the Intrusion Log until the alert flag for that port has been reset.*

When a security violation occurs on a port configured for Port Security, the switch responds in the following ways to notify you:

- The switch sets an alert flag for that port. This flag remains set until:
 - You use either the CLI, menu interface, or web browser interface to reset the flag.
 - The switch is reset to its factory default configuration.

- The switch enables notification of the intrusion through the following means:
 - In the CLI:
 - The **show port-security intrusion-log** command displays the Intrusion Log
 - The **log** command displays the Event Log
 - In the menu interface:
 - The Port Status screen includes a per-port intrusion alert
 - The Event Log includes per-port entries for security violations
 - In the web browser interface:
 - The Alert Log's Status | Overview window includes entries for per-port security violations
 - The Intrusion Log in the Security | Intrusion Log window lists per-port security violation entries
 - In network management applications such as ProCurve Manager via an SNMP trap sent to a network management station

How the Intrusion Log Operates

When the switch detects an intrusion attempt on a port, it enters a record of this event in the Intrusion Log. No further intrusion attempts on that port will appear in the Log until you acknowledge the earlier intrusion event by resetting the alert flag.

The Intrusion Log lists the 20 most recently detected security violation attempts, regardless of whether the alert flags for these attempts have been reset. This gives you a history of past intrusion attempts. Thus, for example, if there is an intrusion alert for port A1 and the Intrusion Log shows two or more entries for port 1, only the most recent entry has not been acknowledged (by resetting the alert flag). The other entries give you a history of past intrusions detected on port A1.

Status and Counters - Intrusion Log		
Port	MAC Address	Date / Time

A1	080009-e93d4f	03/07/06 21:09:34
A1	080009-e93d4f	03/07/06 10:18:43

Figure 11-11. Example of Multiple Intrusion Log Entries for the Same Port

The log shows the most recent intrusion at the top of the listing. You cannot delete Intrusion Log entries (unless you reset the switch to its factory-default configuration). Instead, if the log is filled when the switch detects a new intrusion, the oldest entry is dropped off the listing and the newest entry appears at the top of the listing.

Keeping the Intrusion Log Current by Resetting Alert Flags

When a violation occurs on a port, an alert flag is set for that port and the violation is entered in the Intrusion Log. The switch can detect and handle subsequent intrusions on that port, but will not log another intrusion on the port until you reset the alert flag for either all ports or for the individual port.

Note on Send-Disable Operation

On a given port, if the intrusion action is to send an SNMP trap and then disable the port (**send-disable**), and an intruder is detected on the port, then the switch sends an SNMP trap, sets the port's alert flag, and disables the port. If you re-enable the port without resetting the port's alert flag, then the port operates as follows:

- The port comes up and will block traffic from unauthorized devices it detects.
- If the port detects another intruder, it will send another SNMP trap, but will not become disabled again unless you first reset the port's intrusion flag.

This operation enables the port to continue passing traffic for authorized devices while you take the time to locate and eliminate the intruder. Otherwise, the presence of an intruder could cause the switch to repeatedly disable the port.

Menu: Checking for Intrusions, Listing Intrusion Alerts, and
Resetting Alert Flags

The menu interface indicates per-port intrusions in the Port Status screen, and provides details and the reset function in the Intrusion Log screen.

1. From the Main Menu select:

3. Status and Counters

4. Port Status

The Intrusion Alert column shows “Yes” for any port on which a security violation has been

----- CONSOLE - MANAGER MODE -----						
Status and Counters - Port Status						
Port	Type	Intrusion Alert	Enabled	Status	Mode	Flow Ctrl
A1	10/100TX	No	Yes	Up	Auto	off
A2	10/100TX	No	Yes	Up	Auto	off
A3	10/100TX	Yes	Yes	Up	Auto	off
A4	10/100TX	No	Yes	Up	Auto	off
A5	10/100TX	No	Yes	Up	Auto	off
A6	10/100TX	No	Yes	Down	Auto	off
A7	10/100TX	No	Yes	Up	Auto	off
A8	10/100TX	No	Yes	Down	Auto	off
Actions-> Back Intrusion log Help						
Return to previous screen.						
Use up/down arrow keys to scroll to other entries, left/right arrow keys to change action selection, and <Enter> to execute action.						

Figure 11-12. Example of Port Status Screen with Intrusion Alert on Port A3

2. Type [I] (Intrusion log) to display the Intrusion Log.

MAC Address of Intruding Device on

----- CONSOLE - MANAGER MODE -----			
Status and Counters - Intrusion Log			
Port	MAC Address	Date / Time	System Time of Intrusion on Port
A3	080009-6563e2	08/08/02 16:58:02	
A1	0060b0-896e00	08/08/02 15:28:21	
A3	080009-cf558f	prior to 08/08/02 10:28:58	
Actions-> Back Reset alert flags Help			
Return to previous screen.			
Use up/down arrow keys to scroll to other entries, left/right arrow keys to change action selection, and <Enter> to execute action.			

System Time of Intrusion on Port

Indicates this intrusion on port A3 occurred prior to a reset (reboot) at the indicated time

Figure 11-13. Example of the Intrusion Log Display

The example in Figure 7-11 shows two intrusions for port A3 and one intrusion for port A1. In this case, only the most recent intrusion at port A3 has not been acknowledged (reset). This is indicated by the following:

- Because the Port Status screen (figure 11-12 on page 11-33) does not indicate an intrusion for port A1, the alert flag for the intrusion on port A1 has already been reset.
- Since the switch can show only one uncleared intrusion per port, the alert flag for the older intrusion for port A3 in this example has also been previously reset.

(The intrusion log holds up to 20 intrusion records and deletes an intrusion record only when the log becomes full and a new intrusion is subsequently detected.)

Note also that the “**prior to**” text in the record for the earliest intrusion means that a switch reset occurred at the indicated time and that the intrusion occurred prior to the reset.

3. To acknowledge the most recent intrusion entry on port A3 and enable the switch to enter a subsequently detected intrusion on this port, type **[R]** (for **Reset alert flags**). (Note that if there are unacknowledged intrusions on two or more ports, this step resets the alert flags for all such ports.)

If you then re-display the port status screen, you will see that the Intrusion Alert entry for port A3 has changed to “**No**”. That is, your evidence that the Intrusion Alert flag has been acknowledged (reset) is that the Intrusion Alert column in the port status display no longer shows “**Yes**” for the port on which the intrusion occurred (port A3 in this example). (Because the Intrusion Log provides a history of the last 20 intrusions detected by the switch, resetting the alert flags does not change its content. Thus, displaying the Intrusion Log again will result in the same display as in figure 11-13, above.)

CLI: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags

The following commands display port status, including whether there are intrusion alerts for any port(s), list the last 20 intrusions, and either reset the alert flag on all ports or for a specific port for which an intrusion was detected. (The record of the intrusion remains in the log. For more information, refer to “Operating Notes for Port Security” on page 11-38.)

Syntax: show interfaces brief

List intrusion alert status (and other port status information)’.

show port-security intrusion-log

List intrusion log content.

clear intrusion-flags

Clear intrusion flags on all ports.

port-security [e] < port-number > clear-intrusion-flag

Clear the intrusion flag on one or more specific ports.

In the following example, executing **show interfaces brief** lists the switch's port status, which indicates an intrusion alert on port A1.

ProCurve# show interfaces brief

Status and Counters - Port Status

Port	Type	Intrusion Alert	Enabled	Status	Mode	Flow Ctrl
A1	10/100TX	Yes	Yes	Up	10HDx	off
A2	10/100TX	No	Yes	Up	10HDx	off
A3	10/100TX	No	Yes	Up	10HDx	off
A4	10/100TX	No	Yes	Up	10HDx	off

Intrusion Alert on port

Figure 11-14.Example of an Unacknowledged Intrusion Alert in a Port Status Display

If you wanted to see the details of the intrusion, you would then enter the **show port-security intrusion-log** command. For example:

ProCurve# show port-security intrusion-log

Status and Counters - Intrusion Log

Port	MAC Address	Date / Time
A1	080009-e93d4f	07/03/06 21:09:34
A1	080009-21ae84	07/03/06 17:26:27
A1	080009-e93d4f	prior to 07/03/06 17:18:43

MAC Address of latest Intruder on Port A1

Earlier intrusions on port A1 that have already been cleared (that is, the Alert Flag has been reset at least twice before the most recent intrusion

Dates and Times of Intrusions

0 secs

0 secs

Figure 11-15.Example of the Intrusion Log with Multiple Entries for the Same Port

The above example shows three intrusions for port A1. Since the switch can show only one uncleared intrusion per port, the older two intrusions in this example have already been cleared by earlier use of the **clear intrusion-log** or the **port-security < port-list > clear-intrusion-flag** command. (The intrusion log holds up to 20 intrusion records, and deletes intrusion records only when the log becomes full and new intrusions are subsequently added.) The “**prior to**” text in the record for the third intrusion means that a switch reset occurred at the indicated time and that the intrusion occurred prior to the reset.

To clear the intrusion from port A1 and enable the switch to enter any subsequent intrusion for port A1 in the Intrusion Log, execute the port-security **clear-intrusion-flag** command. If you then re-display the port status screen, you will see that the Intrusion Alert entry for port A1 has changed to **“No”**. (Executing **show port-security intrusion-log** again will result in the same display as above, and does not include the Intrusion Alert status.)

```
ProCurve(config)# port-security a1 clear-intrusion-flag
ProCurve(config)# show interfaces brief
```

Intrusion Alert on port A1 is now

Status and Counters - Port Status							
Port	Type	Intrusion Alert	Enabled	Status	Mode	Flow Ctrl	Bcast Limit
A1	10/100TX	No	Yes	Up	10HDx	off	0
A2	10/100TX	No	Yes	Up	10HDx	off	0
A3	10/100TX	No	Yes	Up	10HDx	off	0

Figure 11-16.Example of Port Status Screen After Alert Flags Reset

For more on clearing intrusions, see “Note on Send-Disable Operation” on page 11-32

Using the Event Log To Find Intrusion Alerts

The Event Log lists port security intrusions as:

```
W MM/DD/YY HH:MM:SS FFI: port A3 - Security Violation
```

where **“W”** is the severity level of the log entry and **FFI** is the system module that generated the entry. For further information, display the Intrusion Log, as shown below.

From the CLI. Type the **log** command from the Manager or Configuration level.

Syntax: log < search-text >

For < **search-text** >, you can use **ffi**, **security**, or **violation**. For example:

```
ProCurve(config)# log security
Keys:   W=Warning   I=Information
        M=Major     D=Debug
----  Event Log listing: Events Since Boot  ----
W 08/01/02 01:18:15 FFI: port A2 - Security Violation
W 08/01/02 04:28:08 FFI: port A1 - Security Violation
----  Bottom of Log : Events Listed = 2  ----

ProCurve(config)# log security
Keys:   W=Warning   I=Information
        M=Major     D=Debug
----  Event Log listing: Events Since Boot  ----
----  Bottom of Log : Events Listed = 0  ----
```

Figure 11-17.Example of Log Listing With and Without Detected Security Violations

From the Menu Interface: In the Main Menu, click on **4. Event Log** and use **Next page** and **Prev page** to review the Event Log contents.

For More Event Log Information. See “Using the Event Log To Identify Problem Sources” in the “Troubleshooting” chapter of the *Management and Configuration Guide* for your switch.

Web: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags

1. Check the Alert Log by clicking on the **Status** tab and the **[Overview]** button. If there is a “Security Violation” entry, do the following:
 - a. Click on the **Security** tab.
 - b. Click on **[Intrusion Log]**. “Ports with Intrusion Flag” indicates any ports for which the alert flag has not been cleared.
 - c. To clear the current alert flags, click on **[Reset Alert Flags]**.

To access the web-based Help provided for the switch, click on **[?]** in the web browser screen.

Operating Notes for Port Security

Identifying the IP Address of an Intruder. The Intrusion Log lists detected intruders by MAC address. If you are using ProCurve Manager to manage your network, you can use the device properties page to link MAC addresses to their corresponding IP addresses.

Proxy Web Servers. If you are using the switch's web browser interface through a switch port configured for Static port security, and your browser access is through a proxy web server, then it is necessary to do the following:

- Enter your PC or workstation MAC address in the port's Authorized Addresses list.
- Enter your PC or workstation's IP address in the switch's IP Authorized Managers list. See "Using Authorized IP Managers" in the *Management and Configuration Guide* for your switch.)

Without both of the above configured, the switch detects only the proxy server's MAC address, and not your PC or workstation MAC address, and interprets your connection as unauthorized.

"Prior To" Entries in the Intrusion Log. If you reset the switch (using the Reset button, Device Reset, or Reboot Switch), the Intrusion Log will list the time of all currently logged intrusions as "prior to" the time of the reset.

Alert Flag Status for Entries Forced Off of the Intrusion Log. If the Intrusion Log is full of entries for which the alert flags have not been reset, a new intrusion will cause the oldest entry to drop off the list, but will not change the alert flag status for the port referenced in the dropped entry. This means that, even if an entry is forced off of the Intrusion Log, no new intrusions can be logged on the port referenced in that entry until you reset the alert flags.

LACP Not Available on Ports Configured for Port Security. To maintain security, LACP is not allowed on ports configured for port security. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port(s), and enables port security on that port. For example:

```
ProCurve(config)# port-security e a17 learn-mode static  
address-limit 2  
LACP has been disabled on secured port(s).  
ProCurve(config)#
```

The switch will not allow you to configure LACP on a port on which port security is enabled. For example:

```
ProCurve(config)# int e a17 lacp passive  
Error configuring port A17: LACP and port security cannot  
be run together.  
ProCurve(config)#
```

To restore LACP to the port, you must remove port security and re-enable LACP active or passive.

Using Authorized IP Managers

Contents

Overview	12-2
Options	12-3
Access Levels	12-3
Defining Authorized Management Stations	12-4
Overview of IP Mask Operation	12-4
Menu: Viewing and Configuring IP Authorized Managers	12-5
CLI: Viewing and Configuring Authorized IP Managers	12-6
Listing the Switch's Current Authorized IP Manager(s)	12-6
Configuring IP Authorized Managers for the Switch	12-7
Web: Configuring IP Authorized Managers	12-9
Web Proxy Servers	12-9
How to Eliminate the Web Proxy Server	12-9
Using a Web Proxy Server to Access the Web Browser Interface ...	12-10
Web-Based Help	12-10
Building IP Masks	12-10
Configuring One Station Per Authorized Manager IP Entry	12-10
Configuring Multiple Stations Per Authorized Manager IP Entry ..	12-11
Additional Examples for Authorizing Multiple Stations	12-13
Operating Notes	12-13

Overview

Authorized IP Manager Features

Feature	Default	Menu	CLI	Web
Listing (Showing) Authorized Managers	n/a	page 12-5	page 12-6	page 12-9
Configuring Authorized IP Managers	None	page 12-5	page 12-6	page 12-9
Building IP Masks	n/a	page 12-10	page 12-10	page 12-10
Operating and Troubleshooting Notes	n/a	page 12-13	page 12-13	page 12-13

The Authorized IP Managers feature uses IP addresses and masks to determine which stations (PCs or workstations) can access the switch through the network. This covers access through the following means:

- Telnet and other terminal emulation applications
- The switch's web browser interface
- SNMP (with a correct community name)

Also, when configured in the switch, the Authorized IP Managers feature takes precedence over local passwords, TACACS+, RADIUS, Port-Based Access Control (802.1X), and Port Security. This means that the IP address of a networked management device must be authorized before the switch will attempt to authenticate the device by invoking any other access security features. If the Authorized IP Managers feature disallows access to the device, then access is denied. Thus, with authorized IP managers configured, having the correct passwords is not sufficient for accessing the switch through the network unless the station attempting access is also included in the switch's Authorized IP Managers configuration.

You can use Authorized IP Managers along with other access security features to provide a more comprehensive security fabric than if you use only one or two security options.

Options

You can configure:

- Up to 10 authorized manager *addresses*, where each address applies to either a single management station or a group of stations
- Manager or Operator access privileges (for Telnet, SNMPv1, and SNMPv2c access only)

Caution

Configuring Authorized IP Managers does not protect access to the switch through a modem or direct connection to the Console (RS-232) port. Also, if an unauthorized station “spoofs” an authorized IP address, it can gain management access to the switch even though a duplicate IP address condition exists. For these reasons, you should enhance your network’s security by keeping physical access to the switch restricted to authorized personnel, using the username/password and other security features available in the switch, and preventing unauthorized access to data on your management stations.

Access Levels

Note

The Authorized IP Manager feature can assign an access level to stations using Telnet, SNMPv1, or SNMPv2c for switch access. The access level the switch allows for authorized stations using SSH, SNMPv3, or the web browser interface is determined by the access application itself, and not by the Authorized IP Manager feature.

For each authorized manager address using Telnet, SNMPv1, or SNMPv2c, you can configure either of these access levels:

- **Manager:** Enables full access to all web browser and console interface screens for viewing, configuration, and all other operations available in these interfaces.
 - **Operator:** Allows read-only access from the web browser and console interfaces. (This is the same access that is allowed by the switch’s operator-level password feature.)
-

Defining Authorized Management Stations

- **Authorizing Single Stations:** The table entry authorizes a single management station to have IP access to the switch. To use this method, just enter the IP address of an authorized management station in the Authorized Manager IP column, and leave the IP Mask set to **255.255.255.255**. This is the easiest way to use the Authorized Managers feature. (For more on this topic, see “Configuring One Station Per Authorized Manager IP Entry” on page 12-10.)
- **Authorizing Multiple Stations:** The table entry uses the IP Mask to authorize access to the switch from a defined group of stations. This is useful if you want to easily authorize several stations to have access to the switch without having to type in an entry for every station. All stations in the group defined by the one Authorized Manager IP table entry and its associated IP mask will have the same access level—Manager or Operator. (For more on this topic, refer to “Configuring Multiple Stations Per Authorized Manager IP Entry” on page 12-11.)

To configure the switch for authorized manager access, enter the appropriate *Authorized Manager IP* value, specify an *IP Mask*, and select either **Manager** or **Operator** for the *Access Level*. The IP Mask determines how the Authorized Manager IP value is used to allow or deny access to the switch by a management station.

Overview of IP Mask Operation

The default IP Mask is 255.255.255.255 and allows switch access only to a station having an IP address that is identical to the Authorized Manager IP parameter value. (“255” in an octet of the mask means that only the exact value in the corresponding octet of the Authorized Manager IP parameter is allowed in the IP address of an authorized management station.) However, you can alter the mask and the Authorized Manager IP parameter to specify ranges of authorized IP addresses. For example, a mask of **255.255.255.0** and any value for the Authorized Manager IP parameter allows a range of 0 through 255 in the 4th octet of the authorized IP address, which enables a block of up to 254 IP addresses for IP management access (excluding 0 for the network and 255 for broadcasts). A mask of **255.255.255.252** uses the 4th octet of a given Autho-

Authorized Manager IP address to authorize four IP addresses for management station access. The details on how to use IP masks are provided under “Building IP Masks” on page 12-10.

Note

The IP Mask is a method for recognizing whether a given IP address is authorized for management access to the switch. This mask serves a different purpose than IP subnet masks and is applied in a different manner.

Menu: Viewing and Configuring IP Authorized Managers

From the console Main Menu, select:

2. Switch Configuration ...

7. IP Authorized Managers

Authorized Manager IP	IP Mask	Access Level
13.28.227.101	255.255.255.252	Manager
13.28.227.104	255.255.255.254	Manager
13.28.227.106	255.255.255.0	Operator
13.28.227.125	255.255.255.255	Manager

Actions-> Back **Add** Edit Delete Help

Return to previous screen.
Use up/down arrow keys to change record selection, left/right arrow keys to change action selection, and <Enter> to execute action.

Figure 12-1. Example of How To Add an Authorized Manager Entry

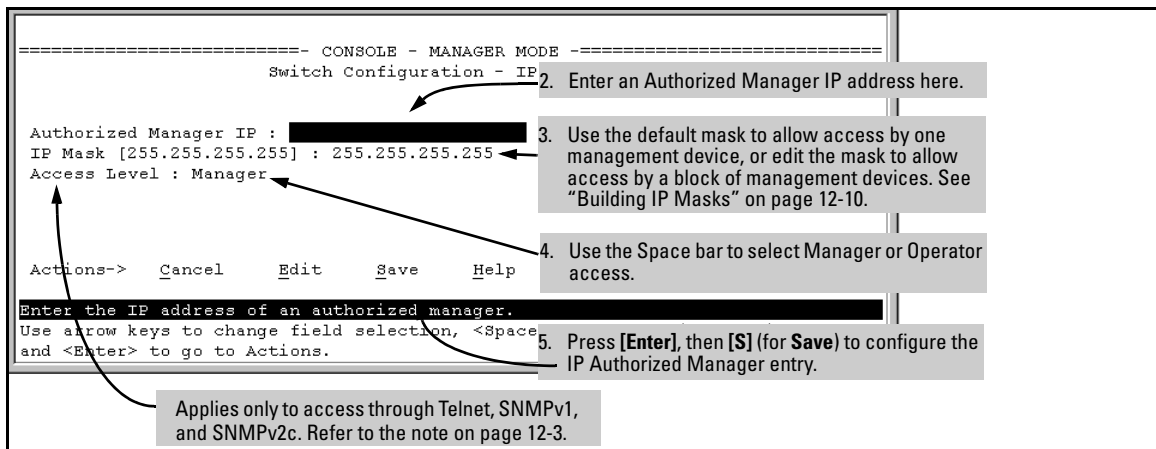


Figure 12-2. Example of How To Add an Authorized Manager Entry (Continued)

Editing or Deleting an Authorized Manager Entry. Go to the IP Managers List screen (figure 12-1), highlight the desired entry, and press **[E]** (for **Edit**) or **[D]** (for **Delete**).

CLI: Viewing and Configuring Authorized IP Managers

Authorized IP Managers Commands Used in This Section

Command	Page
show ip authorized-managers	below
ip authorized-managers	12-7
<ip-address>	12-8
mask <mask-bits>	12-8
<operator manager>	

Listing the Switch’s Current Authorized IP Manager(s)

Use the **show ip authorized-managers** command to list IP stations authorized to access the switch. For example:


```
ProCurve# show ip authorized-managers
IP Managers
Authorized Manager IP   IP Mask               Access Level
-----
10.28.227.101          255.255.255.252      Manager
10.28.227.104          255.255.255.254      Manager
10.28.227.125          255.255.255.255      Manager
10.28.227.106          255.255.255.0        Operator
```

Figure 12-3.Example of the Show IP Authorized-Manager Display

The above example shows an Authorized IP Manager List that allows stations to access the switch as shown below:

IP Mask	Authorized Station IP Address:	Access Mode:
255.255.255.252	10.28.227.100 through 103	Manager
255.255.255.254	10.28.227.104 through 105	Manager
255.255.255.255	10.28.227.125	Manager
255.255.255.0	10.28.227.0 through 255	Operator

Configuring IP Authorized Managers for the Switch

Syntax: ip authorized-managers <ip address>

Configures one or more authorized IP addresses.

[<ip-mask-bits>]

Configures the IP mask for <ip address>

[access <operator | manager>]

Configures the privilege level for <ip address>. Applies only to access through Telnet, SNMPv1, and SNMPv2c. Refer to the Note on page 12-3.

To Authorize Manager Access. This command authorizes manager-level access for any station with an IP address of 10.28.227.0 through 10.28.227.255:

```
ProCurve(config)# ip authorized-managers 10.28.227.101
255.255.255.0 access manager
```

Similarly, the next command authorizes manager-level access for any station having an IP address of 10.28.227.101 through 103:

```
ProCurve(config)# ip authorized-managers 10.28.227.101
255.255.255.252 access manager
```

If you omit the *<mask bits>* when adding a new authorized manager, the switch automatically uses **255.255.255.255**. If you do not specify either Manager or Operator access, the switch assigns the Manager access. For example:

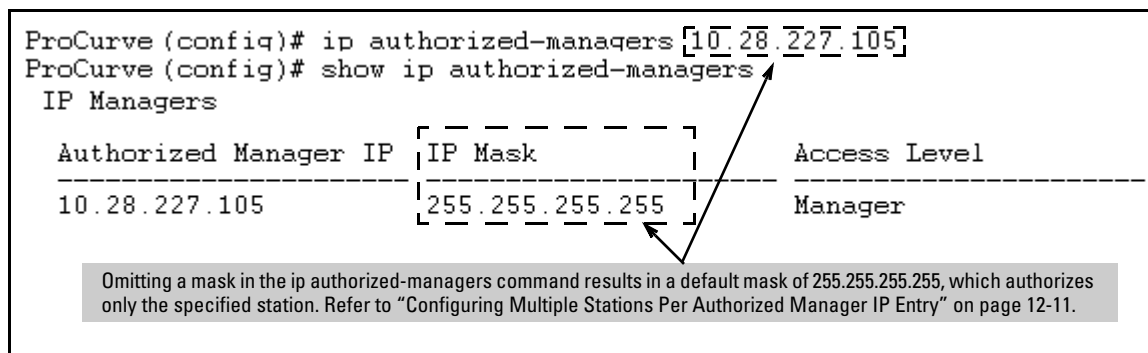


Figure 12-4. Example of Specifying an IP Authorized Manager with the Default Mask

To Edit an Existing Manager Access Entry. To change the mask or access level for an existing entry, use the entry's IP address and enter the new value(s). (Notice that any parameters not included in the command will be set to their default.):

```
ProCurve(config)# ip authorized-managers  
10.28.227.101 255.255.255.0 access operator
```

The above command replaces the existing mask and access level for IP address 10.28.227.101 with 255.255.255.0 and operator.

The following command replaces the existing mask and access level for IP address 10.28.227.101 with 255.255.255.255 and manager (the defaults) because the command does not specify either of these parameters.

```
ProCurve(config)# ip authorized-managers 10.28.227.101
```

To Delete an Authorized Manager Entry. This command uses the IP address of the authorized manager you want to delete:

```
ProCurve(config)# no ip authorized-managers 10.28.227.101
```

Web: Configuring IP Authorized Managers

In the web browser interface you can configure IP Authorized Managers as described below.

To Add, Modify, or Delete an IP Authorized Manager address:

1. Click on the **Security** tab.
2. Click on [Authorized Addresses].
3. Enter the appropriate parameter settings for the operation you want.
4. Click on [Add], [Replace], or [Delete] to implement the configuration change.

Web Proxy Servers

If you use the web browser interface to access the switch from an authorized IP manager station, it is highly recommended that you avoid using a web proxy server in the path between the station and the switch. This is because switch access through a web proxy server requires that you first add the web proxy server to the Authorized Manager IP list. This reduces security by opening switch access to anyone who uses the web proxy server.

How to Eliminate the Web Proxy Server

There are two ways to eliminate a web proxy server from the path between a station and the switch:

1. Add the IP address or DNS name of the switch to the non-proxy or Exceptions list in the web browser interface used on the authorized station.
2. If you don't need proxy server access on the authorized station, disable the proxy server feature in the station's web browser interface.

Note

IP or MAC authentication can be used without a web proxy server.

Using a Web Proxy Server to Access the Web Browser Interface

Caution

This is NOT recommended. Using a web proxy server between the stations and the switch poses a security risk. If the station uses a web proxy server to connect to the switch, any proxy user can access the switch.

If it is necessary to use the switch's web browser interface and your browser access is through a web proxy server, perform these steps:

1. Enter the web proxy server's MAC address in the port's Authorized Addresses list.
2. Enter the web proxy server's IP address in the switch's IP Authorized Managers list.

You must perform both of these steps or the switch only detects the proxy server's MAC address and IP address instead of your workstation addresses, and your connection is considered unauthorized.

Web-Based Help

For web-based help on how to use the web browser interface screen, click on the [?] button provided on the web browser screen.

Building IP Masks

The IP Mask parameter controls how the switch uses an Authorized Manager IP value to recognize the IP addresses of authorized manager stations on your network.

Configuring One Station Per Authorized Manager IP Entry

This is the easiest way to apply a mask. If you have ten or fewer management and/or operator stations, you can configure them by adding the address of each to the Authorized Manager IP list with **255.255.255.255** for the corresponding mask. For example, as shown in figure 12-3 on page 12-7, if you configure an IP address of **10.28.227.125** with an IP mask of **255.255.255.255**, only a station having an IP address of **10.28.227.125** has management access to the switch.

Figure 12-5. Analysis of IP Mask for Single-Station Entries

	1st Octet	2nd Octet	3rd Octet	4th Octet	Manager-Level or Operator-Level Device Access
IP Mask	255	255	255	255	The “255” in each octet of the mask specifies that only the exact value in that octet of the corresponding IP address is allowed. This mask allows management access only to a station having an IP address of 10.33.248.5.
Authorized Manager IP	10	28	227	125	

Configuring Multiple Stations Per Authorized Manager IP Entry

















The mask determines whether the IP address of a station on the network meets the criteria you specify. That is, for a given Authorized Manager entry, the switch applies the IP mask to the IP address you specify to determine a range of authorized IP addresses for management access. As described above, that range can be as small as one IP address (if **255** is set for all octets in the mask), or can include multiple IP addresses (if one or more octets in the mask are set to less than **255**).

If a bit in an octet of the mask is “on” (set to 1), then the corresponding bit in the IP address of a potentially authorized station must match the same bit in the IP address you entered in the Authorized Manager IP list. Conversely, if a bit in an octet of the mask is “off” (set to 0), then the corresponding bit in the IP address of a potentially authorized station on the network does not have to match its counterpart in the IP address you entered in the Authorized Manager IP list. Thus, in the example shown above, a “255” in an IP Mask octet (*all* bits in the octet are “on”) means only one value is allowed for that octet—the value you specify in the corresponding octet of the Authorized Manager IP list. A “0” (all bits in the octet are “off”) means that any value from 0 to 255 is allowed in the corresponding octet in the IP address of an authorized station. You can also specify a series of values that are a subset of the 0-255 range by using a value that is greater than 0, but less than 255.

Figure 12-6. Analysis of IP Mask for Multiple-Station Entries

	1st Octet	2nd Octet	3rd Octet	4th Octet	Manager-Level or Operator-Level Device Access
IP Mask	255	255	255	0	The "255" in the first three octets of the mask specify that only the exact value in the octet of the corresponding IP address is allowed. However, the zero (0) in the 4th octet of the mask allows any value between 0 and 255 in that octet of the corresponding IP address. This mask allows switch access to any device having an IP address of 10.28.227.xxx, where xxx is any value from 0 to 255.
Authorized Manager IP	10	28	227	125	
IP Mask	255	255	255	249	In this example (figure 12-7, below), the IP mask allows a group of up to 4 management stations to access the switch. This is useful if the only devices in the IP address group allowed by the mask are management stations. The "249" in the 4th octet means that bits 0 and 3 - 7 of the 4th octet are fixed. Conversely, bits 1 and 2 of the 4th octet are variable. Any value that matches the authorized IP address settings for the fixed bits is allowed for the purposes of IP management station access to the switch. Thus, any management station having an IP address of 10.28.227. <u>121</u> , <u>123</u> , <u>125</u> , or <u>127</u> can access the switch.
Authorized IP Address	10	28	227	125	

Figure 12-7. Example of How the Bitmap in the IP Mask Defines Authorized Manager Addresses

4th Octet of IP Mask:		249						
4th Octet of Authorized IP Address:		5						
Bit Numbers	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
Bit Values	128	64	32	16	8	4	2	1
4th Octet of IP Mask (249)								
4th Octet of IP Authorized Address (125)								
<p>Bits 1 and 2 in the mask are "off", and bits 0 and 3 - 7 are "on", creating a value of 249 in the 4th octet. Where a mask bit is "on", the corresponding bit setting in the address of a potentially authorized station must match the IP Authorized Address setting for that same bit. Where a mask bit is "off" the corresponding bit setting in the address can be either "on" or "off". In this example, in order for a station to be authorized to access the switch:</p> <ul style="list-style-type: none"> • The first three octets of the station's IP address must match the Authorized IP Address. • Bit 0 and Bits 3 through 6 of the 4th octet in the station's address must be "on" (value = 1). • Bit 7 of the 4th octet in the station's address must be "off" (value = 0). • Bits 1 and 2 can be either "on" or "off". <p>This means that stations with the IP address 13.28.227.X (where X is 121, 123, 125, or 127) are authorized.</p>								

Additional Examples for Authorizing Multiple Stations

	Entries for Authorized Manager List				Results
IP Mask	255	255	0	255	This combination specifies an authorized IP address of 10.33.xxx.1. It could be applied, for example, to a subnetted network where each subnet is defined by the third octet and includes a management station defined by the value of "1" in the fourth octet of the station's IP address.
Authorized Manager IP	10	33	248	1	
IP Mask	255	238	255	250	Allows 230, 231, 246, and 247 in the 2nd octet, and 194, 195, 198, 199 in the 4th octet.
Authorized Manager IP	10	247	100	195	

Operating Notes

- **Network Security Precautions:** You can enhance your network's security by keeping physical access to the switch restricted to authorized personnel, using the password features built into the switch, using the additional security features described in this manual, and preventing unauthorized access to data on your management stations.
- **Modem and Direct Console Access:** Configuring authorized IP managers does not protect against access to the switch through a modem or direct Console (RS-232) port connection.
- **Duplicate IP Addresses:** If the IP address configured in an authorized management station is also configured (or "spoofed") in another station, the other station can gain management access to the switch even though a duplicate IP address condition exists.
- **Web Proxy Servers:** If you use the web browser interface to access the switch from an authorized IP manager station, it is recommended that you avoid the use of a web proxy server in the path between the station and the switch. This is because switch access through a web proxy server requires that you first add the web proxy server to the Authorized Manager IP list. *This reduces security by opening switch access to anyone who uses the web proxy server.* The following two options outline how to eliminate a web proxy server from the path between a station and the switch:

- Even if you need proxy server access enabled in order to use other applications, you can still eliminate proxy service for web access to the switch. To do so, add the IP address or DNS name of the switch to the non-proxy, or “Exceptions” list in the web browser interface you are using on the authorized station.
- If you don’t need proxy server access at all on the authorized station, then just disable the proxy server feature in the station’s web browser interface.

Index

Numerics

3DES ... 7-3

802.1X access control

- authenticate users ... 10-5
- authentication methods ... 10-4
- authentication, local ... 10-6
- authentication, user-based ... 10-4
- authenticator ... 10-19
 - operation ... 10-9
 - show commands ... 10-51
 - unblock port ... 10-6
- authorized-client VLAN, defined ... 10-6
- auth-vid ... 10-23
- auto ... 10-21
- blocked port, trunked ... 10-13
- caution, unauthorized-client VLAN ... 10-37
- CHAP ... 10-3
- chap-radius ... 10-24
- clear-statistics ... 10-26
- client, effect of disconnect ... 10-37
- client-limit, no ... 10-20
- configure
 - commands ... 10-18
 - displaying configuration ... 10-51
 - overview ... 10-17
 - port ... 10-19
- configuring method ... 10-24
- control all clients ... 10-12
- control command ... 10-21
- convert to port-based ... 10-20
- CoS override ... 10-60
- counters ... 10-51
- delay move to unauthorized-client VLAN ... 10-33
- delay Unauth-Client VLAN ... 10-23
- DHCP server ... 10-38
- display all 802.1X, Web, and MAC authentication configuration ... 3-14
- displaying 802.1X port configuration ... 10-53, 10-55, 10-56, 10-57
- EAP ... 10-3
- EAPOL ... 10-7, 10-55
- eap-radius ... 10-24
- enabling controlled directions ... 10-26
- enabling on ports ... 10-19
- enabling on switch ... 10-25
- event log messages ... 10-73
- features ... 10-3
- force authorized ... 10-21, 10-61
- force unauthorized ... 10-21, 10-61
- general setup ... 10-14
- guest VLAN ... 10-7, 10-8, 10-29, 10-36
- GVRP ... 10-65, 10-67
 - effect ... 10-67, 10-72
- hierarchy of precedence in authentication
 - session ... 1-18
- initialize ... 10-26
- LACP not allowed ... 10-73
- local ... 10-24
- logoff-period ... 10-23
- max-requests ... 10-22
- MD5 ... 10-7, 10-48
- meshing, not supported ... 10-13
- messages ... 10-73
- multiple clients ... 10-37
- multiple clients, same VLAN ... 10-5
- open port ... 10-4
- open VLAN
 - authorized client ... 10-31
 - configuration ... 10-41, 10-43
 - general operation ... 10-29
 - mode ... 10-29, 10-35
 - operating notes ... 10-44
 - operating rules ... 10-36
 - PVID, no ... 10-60
 - security breach ... 10-44
 - set up ... 10-40
 - status ... 10-54, 10-61
 - status, viewing ... 10-60
 - suspended VLAN ... 10-61
 - unauthorized client ... 10-31
 - use model ... 10-31
 - VLAN, after authentication ... 10-31, 10-37, 10-44
 - VLAN, tagged ... 10-31, 10-32, 10-37, 10-44, 10-63
- overview ... 10-3

- password for port-access ... 2-11, 2-21
- port, supplicant ... 10-16
- port-based
 - access ... 10-4
 - client without authentication ... 10-5
 - effect of Web/MAC auth operation ... 10-13
 - enable ... 10-19, 10-46
 - latest client, effect ... 10-5
 - multiple client access ... 10-6
 - multiple clients authenticating ... 10-5
 - no client limit ... 10-4
 - open port ... 10-4
 - operation ... 10-5
 - recommended use ... 10-5
 - return to ... 10-20
 - single client authenticates ... 10-5
 - tagged VLAN membership ... 10-5
 - unauthorized client risk ... 10-6
 - untagged VLAN membership ... 10-5, 10-30
 - with Web/MAC authentication ... 10-6
 - See also* user-based.
- port-security ... 10-42
- port-security use ... 10-6
- port-security, with 802.1X ... 10-45
- priority of VLAN, per-port ... 10-10, 10-30
- PVID ... 10-62
- quiet-period ... 10-22
- RADIUS ... 10-3
 - effect on VLAN operation ... 10-65, 10-66
 - host IP address ... 10-25
 - VLAN assignment ... 10-36
- rate-limit override ... 10-60
- reauthenticate ... 10-26
- reauth-period ... 10-23
- rules of operation ... 10-12
- security credentials saved to configuration
 - file ... 2-14, 2-21
- server-timeout ... 10-22
- show commands ... 10-51
- show commands, supplicant ... 10-64
- statistics ... 10-51
- supplicant
 - client not using ... 10-34
 - configuring switch port ... 10-49
 - enabling switch port ... 10-49
 - identity option ... 10-49
 - secret ... 10-49
 - switch port operating as ... 10-47
 - supplicant state ... 10-64
 - supplicant statistics, note ... 10-64
 - supplicant, configuring ... 10-47
 - supplicant-timeout ... 10-22
 - terminology ... 10-6
 - traffic flow on unauthenticated ports ... 10-27
 - troubleshooting, gvrp ... 10-65, 10-66, 10-67
 - trunked port blocked ... 10-13
 - tx-period ... 10-22
 - unauthenticated port ... 10-26
 - unauthorized ... 10-21
 - unauthorized-Client VLAN ... 10-23
 - unauthorized-client VLAN, defined ... 10-8
 - unauthorized-Client VLAN, multiple
 - clients ... 10-39
 - unauth-period ... 10-23
 - unauth-period command ... 10-33
 - unauth-vid ... 10-23
 - use model, open VLAN mode ... 10-31
 - used with port-security ... 10-45
- user-based
 - access ... 10-4
 - authentication ... 10-10
 - client authentication ... 10-5
 - client limit ... 10-3, 10-4, 10-46
 - client-limit, enable ... 10-20
 - clients use same VLAN ... 10-30
 - convert to port-based ... 10-20
 - enable ... 10-19, 10-46
 - limit ... 10-4
 - limit for web auth, MAC auth ... 10-20
 - tagged VLAN ... 10-5
 - VLAN ... 10-38, 10-39
 - Web/MAC authenticated clients ... 10-5
 - See also* port-based.
- user-based vs. port-based ... 10-16
- VLAN
 - authorized-client ... 10-35, 10-36, 10-37
 - guest ... 10-36
 - RADIUS assigned, effect ... 10-38
 - RADIUS override ... 10-35
 - RADIUS-assigned ... 10-36
 - tagged ... 10-33, 10-34
 - temporary membership ... 10-36
 - unauthorized-client ... 10-36, 10-37
 - unauthorized-client, best use ... 10-39
 - unauthorized-client, caution ... 10-37
 - unauthorized-client, on different

- ports ... 10-39
- untagged ... 10-30, 10-33, 10-34
- untagged membership ... 10-20
- VLAN operation ... 10-65
- VLAN use, multiple clients ... 10-7
- VLAN, assignment conflict ... 5-34, 10-12
- VLAN, membership priority ... 10-10, 10-30
- VLAN, priority, RADIUS ... 10-34
- VLAN, tagged membership ... 10-34
- Wake-on-LAN traffic ... 10-27
- Web/MAC auth effect ... 10-13

A

aaa

- gvrp-vlans ... 10-71, 10-72
- port-access gvrp-vlans ... 10-71, 10-72

aaa authentication ... 4-8

- authorized ... 10-24
- chap-radius ... 5-11
- peap-mschapv2 ... 5-11
- privilege-mode ... 4-11
- privilege-mode defined ... 4-12
- web browser ... 5-11

aaa port-access

See Web or MAC authentication.

access

- methods, determining ... 5-7

access levels

- authorized IP managers ... 12-3

accounting

See RADIUS.

address

- authorized for port security ... 11-5

alerts

- generating for monitored events ... 8-34

ARP

- adding IP-to-MAC binding ... 8-20, 8-28
- debugging ... 8-23
- displaying statistics ... 8-22
- dynamic ARP protection ... 8-16
- requests ... 8-33
- trusted ports ... 8-18
- validation checks on ARP packets, configuring ... 8-21

authentication

- concurrent Web/MAC ... 3-4

- DCA-applied parameters to non-authenticated client sessions ... 1-18
- display all 802.1X, Web, and MAC configurations ... 3-14
- NIM override ... 1-18
- RADIUS override ... 1-19

authentication failures ... 8-33

authentication session

- client-specific configuration applied with DCA ... 1-17

authorized addresses

- for IP management security ... 12-4
- for port security ... 11-5

authorized IP managers

- access levels ... 12-3
- building IP masks ... 12-10
- configuring in browser interface ... 12-7, 12-9
- configuring in console ... 12-5
- definitions of single and multiple ... 12-4
- effect of duplicate IP addresses ... 12-13
- IP mask for multiple stations ... 12-11
- IP mask for single station ... 12-10
- IP mask operation ... 12-4
- operating notes ... 12-13
- overview ... 12-1
- troubleshooting ... 12-13

authorized server ... 8-5

authorized server address, configuring ... 8-9

authorized, option for authentication ... 5-11, 10-24

autorun

- autorun-key ... 6-11

autorun-key ... 6-11

B

binding database ... 8-12

bindings

- adding IP to MAC ... 8-28
- displaying configuration ... 8-30

BPDU filtering ... 1-8

BPDU protection ... 1-8

C

cert, RSA key ... 6-11

certificate

- CA-signed ... 7-3

- root ... 7-4
- self-signed ... 7-3
- CHAP** ... 5-11
- chap-radius** ... 5-11
- cipher,SSH** ... 6-17
- Clear button**
 - to delete password protection ... 2-7
- configuration**
 - filters ... 9-2
 - port security ... 11-7
 - RADIUS
 - See *RADIUS*.
 - saving security credentials in multiple
 - files ... 2-20
 - SSH
 - See *SSH*.
 - storage of security credentials
- console**
 - authorized IP managers, configuring ... 12-5
- CoS**
 - override ... 10-60
 - priority assignment ... 5-4
- crypto**
 - babble ... 6-12
 - fingerprint ... 6-12

D

- database, snooping** ... 8-5
- debug logging**
 - DHCP snooping ... 8-13
- default configuration and security** ... 1-15
- default settings**
 - 802.1X
 - aaa port-access authenticator control,
 - auto* ... 10-61
 - aaa port-access supplicant ... 10-49–10-50
 - access control, *none* ... 1-6
 - authentication, *disabled* ... 10-3
 - maximum clients on a port, *one* ... 10-66
 - aaa authentication parameters ... 4-12
 - aaa port-access authenticator control ... 10-21–10-23
 - aaa port-access mac-based ... 3-52–3-54
 - aaa port-access web-based ... 3-24–3-27
 - authorized ip managers, *none* ... 1-5
 - client public-key authentication, *disabled* ... 6-2
 - connection-rate filtering, *none* ... 1-8

- DHCP snooping
 - database parameters ... 8-12
 - on VLANs, *disabled* ... 8-7
 - Option 82 remote-id, *MAC address* ... 8-11
 - Option 82 untrusted-policy, *drop* ... 8-10
 - trusted ports, *disabled* ... 8-8
- DHCP snooping, *none* ... 1-8
- dynamic arp protection, *none* ... 1-8
- dynamic IP lockdown, *none* ... 1-8
- factory reset, *enabled* ... 2-28
- front panel security ... 2-3, 2-27
- ICMP rate-limiting, *none* ... 1-8
- instrumentation monitor
 - SNMP traps, *disabled* ... 8-35
 - thresholds and parameters ... 8-35
- instrumentation monitor, *disabled* ... 8-35
- IP mask, *255.255.255.255* ... 12-4
- MAC authentication, *disabled* ... 1-6
- MAC lockdown and lockout, *none* ... 1-7
- manager password, *no password* ... 1-3
- passwords
 - clear password, *enabled* ... 2-27
 - password recovery, *enabled* ... 2-28, 2-33
 - password-clear, *enabled* ... 2-29
 - reset-on-clear, *disabled* ... 2-28
 - username and passwords, *none* ... 2-3
- port security, *none* ... 1-7
- port security, *off* or '*continuous*' ... 11-4
- RADIUS
 - accounting, *none* ... 5-3
 - authentication, *none* ... 5-3
 - authorization, *none* ... 5-3
 - global parameters ... 5-9
 - radius-server dead-time, *0 minutes* ... 5-18
 - radius-server retransmit, *3 seconds* ... 5-19
 - radius-server timeout, *3 seconds* ... 5-18
 - server key, *null* ... 5-9
- RADIUS authentication, *disabled* ... 1-5
- secure management vlan, *disabled* ... 1-5
- security
 - access security and authentication ... 1-3
 - network security ... 1-7
- SNMP access ... 1-15
- SNMP access to the security MIB, *open* ... 5-21
- SNMP, *public, unrestricted* ... 1-5
- source-port filters ... 9-2
- spanning tree
 - bpdu filtering, *none* ... 1-8

- bpdu protection, *none* ... 1-8
- SSH, *disabled* ... 1-4, 6-2
- SSL, *disabled* ... 1-5, 7-2
- TACACS+
 - authentication configuration ... 4-9
 - authentication, *disabled* ... 1-5, 4-2
 - login attempts, 3 ... 4-6
 - tacacs-server-timeout, 5 seconds ... 4-23
- TCP port number for SSH connections, 22 ... 6-18
- TCP port number for SSL connections, 443 ... 7-19
- Telnet access, *enabled* ... 1-4
- traffic filters, *none* ... 9-2
- traffic/security filters, *none* ... 1-7
- UDP destination port for accounting, 1813 ... 5-7
- UDP destination port for authentication, 1812 ... 5-7
- user authentication, *disabled* ... 6-2
- virus throttling, *none* ... 1-8
- Web and MAC authentication ... 3-3–3-54
- Web authentication, *disabled* ... 1-6
- Web-browser access, *enabled* ... 1-4
- denial-of-service**
 - avoid attacks using DHCP snooping ... 8-4
 - monitoring system resources ... 8-33
- DES** ... 7-3
- DHCP Option 82**
 - IP-to-MAC binding database ... 8-20, 8-28
- DHCP protection**
 - See* DHCP snooping.
- DHCP snooping** ... 8-4
 - authorized server ... 8-5
 - binding database ... 8-12
 - changing remote-id ... 8-11
 - configuring authorized server address ... 8-9
 - database ... 8-5
 - denial-of-service attack ... 8-4
 - DHCPACK ... 8-5
 - DHCPDECLINE ... 8-5
 - DHCPNACK ... 8-5
 - DHCPOFFER ... 8-5
 - DHCPRELEASE ... 8-5
 - disable MAC check ... 8-11
 - disabling ... 8-5
 - dropping packets ... 8-5
 - enabling ... 8-5
 - debug logging ... 8-13
 - on trusted ports ... 8-8
 - on VLANs ... 8-6, 8-7
 - IP-to-MAC binding database ... 8-20, 8-28
 - log messages ... 8-14
 - Option 82 ... 8-9
 - option parameter ... 8-6
 - remote-id ... 8-10
 - show configuration ... 8-6
 - stats ... 8-6
 - trust ... 8-6
 - untrusted-policy ... 8-10
 - verify ... 8-6
- documentation**
 - feature matrix ... -xviii
 - latest versions ... -xvii
 - printed in-box publication ... -xvii
 - release notes ... -xvii
- duplicate IP address**
 - effect on authorized IP managers ... 12-13
- dynamic ARP protection**
 - additional validation checks on ARP packets ... 8-21
 - ARP packet debugging ... 8-23
 - displaying ARP statistics ... 8-22
 - enabling ... 8-16
 - IP-to-MAC binding, adding to DHCP database ... 8-20, 8-28
 - trusted ports, configuring ... 8-18
 - verifying configuration ... 8-21
- Dynamic Configuration Arbiter (DCA)**
 - applying settings to non-authenticated clients ... 1-18
 - hierarchy of precedence in authentication sessions ... 1-19
 - overview ... 1-17
- dynamic IP lockdown**
 - debugging ... 8-31
 - DHCP binding database ... 8-25
 - DHCP leases ... 8-25
 - DHCP snooping ... 8-24
 - enabling ... 8-26
 - filtering IP addresses ... 8-25
 - overview ... 8-23
 - spoofing protection ... 8-24
 - verifying configuration ... 8-29
 - VLAN binding ... 8-25

E

Eavesdrop Protection ... 11-4

encryption key

RADIUS ... 2-11, 2-15

TACACS ... 2-11, 2-15

event log

alerts for monitored events ... 8-34

intrusion alerts ... 11-36

F

filetransfer, SSH ... 6-17

filter, source-port

editing ... 9-18

filter indexing ... 9-19

filter type ... 9-8

idx ... 9-8, 9-19

index ... 9-8, 9-19

named ... 9-6

operating rules ... 9-4, 9-6

port-trunk operation ... 9-2, 9-17

show ... 9-8

value ... 9-8

viewing ... 9-8

filtering MAC addresses ... 8-25

filters

overview ... 9-2

source port ... 9-4

source-port filter value ... 9-19

static ... 9-2

types ... 9-3

Framed-IP-Address ... 5-38

G

guest VLAN ... 10-7, 10-8, 10-29

GVRP

dynamic VLANs ... 10-71

effect on client authentication ... 10-65, 10-67

static VLAN not advertised ... 10-68, 10-72

H

hierarchy of precedence, used by DCA ... 1-17

I

IDM

overview ... 1-21

RADIUS-based security classifiers ... 1-21

include credentials command ... 2-11

See also security credentials.

interface

unknown-vlans command ... 10-72

intrusion alarms

entries dropped from log ... 11-38

event log ... 11-36

prior to ... 11-38

Intrusion Log

prior to ... 11-34, 11-35

IP

address count ... 8-33

authorized IP managers ... 12-1

reserved port numbers ... 6-19

IP attribute ... 5-38

IP masks

building ... 12-10

for multiple authorized manager

stations ... 12-11

for single authorized manager station ... 12-10

operation ... 12-4

IP routing

dynamic ARP protection, enabling ... 8-16

validation checks on ARP packets,

configuring ... 8-21

ip source-binding ... 8-29

ip source-lockdown ... 8-26, 8-27

ip source-lockdown bindings ... 8-30, 8-31

IP spoofing

protection against ... 8-24

IP-to-MAC binding ... 8-20, 8-28

ip-version, ipv4or6 ... 6-17

L

LACP

802.1X not allowed ... 10-13, 10-19, 10-73

login attempts, monitoring ... 8-33

M

MAC address

count, monitoring ... 8-33

inconsistent value ... 11-20

moves, monitoring ... 8-33

MAC authentication

- authenticator operation ... 3-6
- blocked traffic ... 3-3
- CHAP
 - defined ... 3-11
 - usage ... 3-3
- client status ... 3-60
- concurrent with Web ... 3-4
- configuration commands ... 3-51
- configuring
 - on the switch ... 3-50
 - switch for RADIUS access ... 3-17
 - the RADIUS server ... 3-16
- general setup ... 3-14
- hierarchy of precedence in authentication
 - session ... 1-18
- LACP not allowed ... 3-13
- overview ... 1-6
- port access ... 10-4
- rules of operation ... 3-12
- show status and configuration ... 3-54
- terminology ... 3-11
- MAC Lockdown** ... 11-3
- MAC Lockout** ... 11-3
- manager password** ... 2-4, 2-6, 2-7
 - recommended ... 4-7
 - saving to configuration file ... 2-12
- MD5**
 - See* RADIUS.
- message**
 - inconsistent value ... 11-20
- MIB**
 - SNMP access ... 1-15
 - SNMP access to authentication MIB ... 1-15
- MSCHAPv2** ... 5-11

N

- named source port filters**
 - configuring ... 9-7
 - displaying ... 9-8
 - operating rules ... 9-6
 - See also* source port filters.
- Network Immunity Manager (NIM)**
 - hierarchy of precedence in authentication
 - session ... 1-18
 - overview ... 1-18

O

- OOBM** ... 3-18, 4-19, 4-20, 5-15, 5-40, 6-9, 6-18, 6-27, 6-28, 10-25, 10-42
- oobm** ... 4-19
- open VLAN mode**
 - See* 802.1X access control.
- OpenSSH** ... 6-3
- OpenSSL** ... 7-2
- operating notes**
 - authorized IP managers ... 12-13
 - port security ... 11-38
- operator password** ... 2-4, 2-6, 2-7
 - saving to configuration file ... 2-12
- Option 82**
 - snooping ... 8-6

P

- packet validation** ... 8-6
- password**
 - 802.1X port-access ... 2-11, 2-21
 - browser/console access ... 2-5
 - case-sensitive ... 2-6
 - caution ... 2-5
 - configuring manager and operator ... 2-12
 - delete ... 2-7
 - deleting with the Clear button ... 2-7
 - if you lose the password ... 2-7
 - included with security credentials ... 2-11
 - incorrect ... 2-5
 - length ... 2-6
 - locally configured, hierarchy of precedence in
 - authentication session ... 1-18
 - operator only, caution ... 2-5
 - password pair ... 2-3
 - setting ... 2-6
 - SNMP configuration ... 2-3
- password security** ... 6-20
 - saved to configuration file ... 2-12
- peap-mschapv2** ... 5-11
- port**
 - security configuration ... 11-3
 - trusted ... 8-18
 - untrusted ... 8-20, 8-28
- port access**
 - client limit ... 10-20
 - concurrent ... 10-20
 - MAC authentication ... 10-4

- tracking client authentication failures ... 8-33
- Web authentication ... 10-4
- Web/MAC ... 10-20
- See also* 802.1X access control.
- port scan, detecting** ... 8-33
- port security**
 - 802.1X, learn mode requirement ... 11-14
 - authorized address definition ... 11-5
 - basic operation ... 11-4
 - caution, device limit ... 11-14
 - configuring ... 11-7
 - configuring in browser interface ... 11-30, 11-37
 - event log ... 11-36
 - notice of security violations ... 11-30
 - operating notes ... 11-38
 - overview ... 1-7, 11-3
 - prior to ... 11-38
 - proxy web server ... 11-38
 - TCP/UDP closed ports ... 8-33
- port-based access control**
 - password ... 2-11, 2-21
 - port-security learn mode ... 11-14
 - VLAN, tagged member ... 10-31
 - See also* 802.1X access control.
- ports**
 - trusted ... 8-6
- prior to** ... 11-34, 11-35, 11-38
- Privacy Enhanced Mode (PEM)**
 - See* SSH.
- privilege-mode** ... 4-11, 4-12
- ProCurve**
 - switch documentation ... -xvii
- ProCurve Manager**
 - IDM as a plug-in to ... 1-21
 - port security alerts ... 11-4
- proxy**
 - web server ... 11-38

R

RADIUS

- accounting ... 5-4, 5-37
- accounting, configuration outline ... 5-39
- accounting, configure server access ... 5-40
- accounting, configure types on switch ... 5-42
- accounting, exec ... 5-38, 5-42
- accounting, interim updating ... 5-44
- accounting, network ... 5-42, 5-43
- accounting, operating rules ... 5-39
- accounting, server failure ... 5-39
- accounting, session-blocking ... 5-44
- accounting, start-stop method ... 5-43
- accounting, statistics terms ... 5-47
- accounting, stop-only method ... 5-43
- accounting, system ... 5-38, 5-42
- administrative-user service-type value ... 5-14
- authentication options ... 5-3
- authentication, authorized ... 5-11
- authentication, local ... 5-24
- authentication, web ... 5-10
- authentication, web browser ... 5-11, 5-14
- authorization ... 5-26
- bypass RADIUS server ... 5-13
- change-of-authorization attribute ... 5-36
- commands accounting ... 5-38
- commands authorization ... 5-26
- commands, accounting ... 5-37
- commands, switch ... 5-8
- configuration outline ... 5-9
- configure server access ... 5-14
- configuring commands authorization ... 5-28
- configuring switch global parameters ... 5-17
- Egress-VLAN ID attribute ... 5-35
- Egress-VLAN-Name attribute ... 5-35
- Framed-IP-Address ... 5-38
- general setup ... 5-7
- HP-acct-terminate-cause attribute ... 5-36
- HP-Command-Exception ... 5-28
- HP-command-string ... 5-28
- IP attribute ... 5-38
- local authentication ... 5-13
- login privilege-mode, application options ... 5-14
- login-privilege mode ... 5-13
- manager access denied ... 5-14
- manager access privilege ... 5-13
- manager access privilege, service type
 - value ... 5-8
- MD5 ... 5-6
- messages ... 5-53
- MS-RAS-Vendor attribute ... 5-36
- NAS-Prompt-User service-type value ... 5-14
- network accounting ... 5-37
- operating rules, switch ... 5-6
- security ... 5-13
- security note ... 5-4
- server access order ... 5-39

- server access order, changing ... 5-50
- servers, multiple ... 5-19
- service type value ... 5-8
- service-type value ... 5-14
- service-type value, null ... 5-14
- shared secret key, saving to configuration file ... 2-11, 2-15
- show accounting ... 5-49
- show authentication ... 5-48
- SNMP access security not supported ... 5-4
- SNMP access to auth config MIB ... 5-4
- statistics, viewing ... 5-46
- terminology ... 5-5
- TLS ... 5-6
- Tunnel-Type attribute ... 5-35
- vendor specific attributes ... 5-36
- vendor-specific attributes ... 5-28
- VSAs ... 5-29
- web browser security not supported ... 5-7
- web-browser access controls ... 5-25
- web-browser security not supported ... 5-4, 5-25
- rate-limit**
 - override ... 10-60
- rate-limiting** ... 8-34
 - RADIUS administered ... 5-4
- remote-id**
 - changing ... 8-11
 - snooping ... 8-10
- reserved port numbers** ... 6-19, 7-20
- Reset-on-clear**
 - disabled when saving security credentials to configuration file ... 2-20
- RFCs**
 - RFC 2548 ... 5-36
 - RFC 3580 ... 5-35
 - RFC 4675 ... 5-35
- RSA key, cert** ... 6-11

S

- security**
 - authorized IP managers ... 12-1
 - per port ... 11-3
- security credentials**
 - 802.1X credentials saved to configuration file ... 2-14, 2-21
 - copying configurations on the switch ... 2-20
 - copying startup configuration ... 2-19

- disabling Reset-on-clear option ... 2-20
- downloading a configuration file ... 2-19
- downloading from a server ... 2-10
- enabling storage in configuration file ... 2-11
- manager username and password ... 2-12
- operator username and password ... 2-12
- RADIUS ... 2-15
- saving to running configuration with include-credentials ... 2-19
- saving to startup configuration ... 2-19
- saving to startup configuration with write memory ... 2-19
- SNMPv3 ... 2-13
- SSH ... 2-16
- SSH private keys not saved ... 2-21
- TACACS ... 2-15
- viewing in running configuration ... 2-11
- viewing in startup configuration ... 2-19
- when SNMPv3 credentials in downloaded file are not supported ... 2-21

security violations

- detecting ... 8-33
- notices of ... 11-30

security, password

See SSH.

setting a password ... 2-6

shared secret key

- RADIUS ... 2-11, 2-15
- TACACS ... 2-11, 2-15

single sign-on ... 4-11

SNMP

- authentication failures ... 8-33
- disabling access to authentication MIB ... 1-15
- password and username configuration ... 2-3
- RADIUS access to auth config MIB ... 5-4
- trap generation ... 8-35

SNMPv3

- saving security credentials to configuration file ... 2-11, 2-13
- security credentials not supported in downloaded file ... 2-21

snooping

- authorized server ... 8-5
- authorized server address ... 8-9
- binding database ... 8-12
- changing remote-id ... 8-11
- DHCP ... 8-4
- disable MAC check ... 8-11

- Option 82 ... 8-6, 8-9
- statistics ... 8-6
- untrusted-policy ... 8-10
- verify ... 8-6
- source port filters**
 - configuring ... 9-4
 - named ... 9-6
 - operating rules ... 9-4
 - See also* named source port filters.
 - selection criteria ... 9-3
- spanning tree**
 - edge port configuration ... 3-22, 10-26
 - security features ... 1-8
- spoofing**
 - protection against ... 8-24
- SSH**
 - authenticating switch to client ... 6-3
 - authentication, client public key ... 6-2
 - authentication, user password ... 6-2
 - caution, security ... 6-19
 - cipher ... 6-17
 - CLI commands ... 6-9
 - client behavior ... 6-15, 6-16
 - client public-key authentication ... 6-21, 6-24
 - client public-key, clearing ... 6-29
 - client public-key, creating file ... 6-26
 - client public-key, displaying ... 6-28
 - configuring authentication ... 6-20
 - configuring key lengths ... 6-13
 - crypto key ... 6-11
 - disabling ... 6-11
 - enable ... 6-16, 7-19
 - enabling ... 6-15
 - erase host key pair ... 6-11
 - filetransfer ... 6-17
 - generate host key pair ... 6-11
 - generating key pairs ... 6-10
 - host key pair ... 6-11
 - ip-version ... 6-17
 - key, babble ... 6-12
 - key, fingerprint ... 6-12
 - keys, zeroing ... 6-11
 - keysize ... 6-13
 - known-host file ... 6-14, 6-15
 - mac selection ... 6-18
 - man-in-the-middle spoofing ... 6-16
 - messages, operating ... 6-30
 - OpenSSH ... 6-3
 - operating rules ... 6-8
 - outbound SSH not secure ... 6-8
 - password security ... 6-20
 - password-only authentication ... 6-20
 - passwords, assigning ... 6-9
 - PEM ... 6-3
 - prerequisites ... 6-5
 - private keys not saved to configuration
 - file ... 2-21
 - public key ... 6-5, 6-14
 - public key, displaying ... 6-14
 - public key, saving to configuration file ... 2-11, 2-16
 - reserved IP port numbers ... 6-19
 - security ... 6-19
 - SSHv2 ... 6-2
 - steps for configuring ... 6-6
 - switch key to client ... 6-13
 - terminology ... 6-3
 - unauthorized access ... 6-30
 - version ... 6-2
 - zeroing a key ... 6-11
 - zeroize ... 6-11
- SSL**
 - CA-signed ... 7-3, 7-15
 - CA-signed certificate ... 7-3, 7-15
 - CLI commands ... 7-7
 - client behavior ... 7-17, 7-18
 - crypto key ... 7-10
 - disabling ... 7-9, 7-17
 - enabling ... 7-17
 - erase certificate key pair ... 7-9
 - erase host key pair ... 7-9
 - generate CA-signed ... 7-15
 - generate CA-signed certificate ... 7-15
 - generate host key pair ... 7-9
 - generate self-signed ... 7-12
 - generate self-signed certificate ... 7-9, 7-12
 - generate server host certificate ... 7-9
 - generating Host Certificate ... 7-8
 - host key pair ... 7-9
 - key, babble ... 7-12
 - key, fingerprint ... 7-12
 - man-in-the-middle spoofing ... 7-18
 - OpenSSL ... 7-2
 - operating notes ... 7-6
 - operating rules ... 7-6
 - passwords, assigning ... 7-7

- prerequisites ... 7-5
- remove self-signed certificate ... 7-9
- remove server host certificate ... 7-9
- reserved TCP port numbers ... 7-20
- root ... 7-4
- root certificate ... 7-4
- self-signed ... 7-3, 7-12
- self-signed certificate ... 7-3, 7-9, 7-12
- server host certificate ... 7-9
- SSL server ... 7-3
- SSLv3 ... 7-2
- steps for configuring ... 7-5
- supported encryption methods ... 7-3
- terminology ... 7-3
- TLSv1 ... 7-2
- troubleshooting, operating ... 7-21
- unsecured web browser access ... 7-18
- version ... 7-2
- zeroize ... 7-10, 7-11

ssl

- cert ... 6-11

static filter limit ... 9-2

STP

- prerequisite for 802.1X controlled directions ... 10-27
- prerequisite for web-based controlled directions ... 3-22

system delay ... 8-33

system resource usage ... 8-33

T

TACACS

- aaa parameters ... 4-12
- authentication ... 4-3
- authentication process ... 4-24
- authentication, local ... 4-26
- authorized IP managers, effect ... 4-29
- configuration, authentication ... 4-11
- configuration, encryption key ... 4-22
- configuration, server access ... 4-18
- configuration, timeout ... 4-23
- configuration, viewing ... 4-10
- encryption key ... 4-6, 4-18, 4-19, 4-22
- encryption key exclusion ... 4-30
- encryption key, general operation ... 4-27
- encryption key, global ... 4-23

- encryption key, saving to configuration file ... 2-11
- general operation ... 4-2
- IP address, server ... 4-18
- local manager password requirement ... 4-30
- messages ... 4-29
- NAS ... 4-3
- precautions ... 4-5
- preparing to configure ... 4-8
- preventing switch lockout ... 4-18
- privilege level code ... 4-7
- server access ... 4-18
- server priority ... 4-21
- setup, general ... 4-5
- shared secret key, saving to configuration file ... 2-15
- show authentication ... 4-8
- single login ... 4-13
- single sign-on ... 4-13
- system requirements ... 4-5
- TACACS+ server ... 4-3, 4-8
- testing ... 4-5
- TFTP, configuration ... 4-30
- timeout ... 4-18
- troubleshooting ... 4-6
- unauthorized access, preventing ... 4-7
- web access, controlling ... 4-28
- web access, no effect on ... 4-5

TCP

- reserved port numbers ... 7-20

TCP/UDP

- monitoring packets to closed ports ... 8-33

test ... 4-18

TLS

- See* RADIUS.

troubleshooting

- authentication via Telnet ... 4-18
- authorized IP managers ... 12-13

trunk

- filter, source-port ... 9-2, 9-17
- LACP, 802.1X not allowed ... 10-19
- See also* LACP.

trusted port

- defined ... 8-18
- enabling ... 8-18

trusted ports ... 8-6

- enabling snooping ... 8-8

U

untrusted policy, snooping ... 8-10

user name

cleared ... 2-7

SNMP configuration ... 2-3

V

vendor-specific attribute

configuring support for HP VSAs ... 5-29

defining ... 5-30

virus detection

monitoring ARP requests ... 8-33

VLAN

802.1X ... 10-65

802.1X, ID changes ... 10-68, 10-72

802.1X, suspend untagged VLAN ... 10-61

not advertised for GVRP ... 10-68, 10-72

Tagged egress VLAN in authentication

session ... 5-35

tagged egress VLAN in authentication

session ... 5-35

untagged VLAN in authentication

session ... 5-34, 5-35

VLANs

GVRP-created ... 10-72

GVRP-learned ... 10-71

See also VLAN.

W

Wake-on-LAN

on 802.1X-aware ports ... 10-27

on web-authenticated ports ... 3-23

warranty ... i

Web authentication

authenticator operation ... 3-6

blocked traffic ... 3-3

CHAP

defined ... 3-11

usage ... 3-3

client status ... 3-60

concurrent with MAC ... 3-4

configuration commands ... 3-21

configuring

access control on unauthenticated

ports ... 3-22

controlled directions ... 3-22

on the switch ... 3-20

switch for RADIUS access ... 3-17

display all 802.1X, Web, and MAC authentication

configuration ... 3-14, 10-15

general setup ... 3-14

hierarchy of precedence in authentication

session ... 1-18

LACP not allowed ... 3-13

overview ... 1-6

port access ... 10-4

redirect URL ... 3-11

rules of operation ... 3-12

show status and configuration ... 3-28

terminology ... 3-11

Wake-on-LAN traffic ... 3-23

web browser interface

aaa authentication ... 5-8

access ... 5-7

authorized IP managers, configuring ... 12-7,
12-9

configuring port security ... 11-30, 11-37

SSL ... 7-18

unsecured access, SSL ... 7-18

web server, proxy ... 11-38

Technology for better business outcomes

To learn more, visit
www.hp.com/go/ bladesystem/documentation/

© Copyright 2009 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP will not be liable for technical or editorial errors or omissions contained herein.



August 2009

Manual Part Number
5992-5525

